arXiv:1412.6697v2 [math.ST] 26 Mar 2015

MODEL SELECTION FOR THE SEGMENTATION OF MULTIPARAMETER EXPONENTIAL FAMILY DISTRIBUTIONS

ALICE CLEYNEN AND EMILIE LEBARBIER

Abstract. We consider the segmentation problem of univariate distributions from the exponential family with multiple parameters. In segmentation, the choice of the number of segments remains a difficult issue due to the discrete nature of the change-points. In this general exponential family distribution framework, we propose a penalized log-likelihood estimator where the penalty is inspired by papers of L. Birgé and P. Massart. The resulting estimator is proved to satisfy an oracle inequality. We then further study the particular case of categorical variables by comparing the values of the key constants when derived from the specification of our general approach and when obtained by working directly with the characteristics of this distribution. Finally, a simulation study is conducted to assess the performance of our criterion for the exponential distribution, and an application on real data modeled by the categorical distribution is provided.

1. Introduction

Penalized-likelihood approaches are becoming more and more popular in numerous domains in statistics: density estimation, variable selection, machine learning, etc. Here we consider a multiple change-point detection setting for univariate datasets where we are interested in the estimation of the number of segments K. While earlier works have focused on variables distributed from specified distributions, for instance Gaussian [1], or either Poisson or negative binomial distributions [2], we consider here a more general framework of exponential family distributions. More precisely, the model is the following:

$$Y_t \sim \mathcal{G}(\boldsymbol{\theta}_t) = s(t),\ 1 \leq t \leq n$$

where $\mathcal{G}$ is a distribution from the exponential family with parameter $\boldsymbol{\theta}_t \in \mathbb{R}^d$. In the context of change-point models, we want to consider that $\boldsymbol{\theta}_t$ is piece-wise constant along the time-line $1, \ldots, n$ and we therefore wish to identify a partition of $\{1, \ldots, n\}$ into K segments within which the observations can be modeled as following the same distribution while they differ between segments.

In our framework we will consider the minimal canonical form of the exponential family distribution which we will write as

$$s(t) = h(Y_t) \exp\left[\boldsymbol{\theta}_t . \mathbf{T}_t - A(\boldsymbol{\theta}_t)\right]$$

where we use the . symbol to denote the canonical scalar product of $\mathbb{R}^d$, A is the log-partition function of $\mathcal{G}$ and $\mathbf{T}_t = T(Y_t)$ $(\in \mathbb{R}^d)$ is the minimal sufficient statistic associated with variable Y_t.

While almost all methods for choosing the number of segments can be seen as penalized-likelihood approaches (Akaike Information Criterion, [3], Bayes Information Criterion [4], Integrated Completed Likelihood [5], etc), we and other authors (see for instance [1, 6, 7]) have previously emphasized how crucial the choice of the penalty function is in contexts such as segmentation where the size of the collection of models grows with the size of the data. For this reason we have extended the approach developed in the Poisson and negative binomial cases to more general distributions from the exponential family. The motivation for this work initially relied on the very strong similarities which we had observed between the Poisson and negative binomial distributions. Most of those similarities are in fact related to properties of the log-partition function A.

1991 *Mathematics Subject Classification.* primary 62G05, 62G07; secondary 62P10.

Key words and phrases. Exponential Family; Distribution estimation; Change-point detection; Model selection.

Exponential family distributions, and in particular the log-partition function, have been well studied in the past years. In a pioneer work [8], Brown has described the fundamental properties of exponential family distributions, such as parametrization using sufficient statistics, differentiability of the log-partition function and its relation to moments, etc. More recently, [9] has demonstrated the strong links between graphical models and exponential family, [10] has studied the sub-exponential growth of the cumulants of an exponential family distribution and studied convergence rates of regularization algorithm under sparsity assumptions while [11] has studied consistency properties of the lasso procedure under some convexity assumption.

In this paper our goal is to mimic the procedure followed in the Poisson and negative binomial cases for obtaining the oracle inequality while emphasizing the role of the log-partition function. Considering the essential role played by the chi-square statistic in this earlier work, we will restrict the considered families to those with positive marginal sufficient statistics (hence allowing the definition of χ^2), thus we will assume that the set of natural parameters $\boldsymbol{\theta}_t$ is restricted to a compact $\boldsymbol{\nu}$ such that $\nabla A(\boldsymbol{\nu})$ belongs to $\{\mathbb{R}^+\}^d$. This for instance excludes the Gaussian distribution unless we assume the mean parameter to be positive.

Among the key features of minimal exponential family is the relationship between the derivatives of A and the moments of the sufficient statistics. Hence the first two moments are given by

- $\mathbb{E}(\mathbf{T}_t) = \nabla A(\boldsymbol{\theta}_t)$
- $Cov(\mathbf{T}_t) = \nabla^2 A(\boldsymbol{\theta}_t)$

Moreover, using minimal representation of the exponential family ensures that the gradient mapping $\nabla A : \boldsymbol{\nu} \to \nabla A(\boldsymbol{\nu})$ is a bijection (see for instance [9]). These properties will be fundamental in the construction of the oracle inequality. The proof of the later mimicks the approach adopted in [2], most of the time studying each marginal of the sufficient statistic separately.

In Section 2 we introduce the collection of models and the penalized-likelihood framework and state our main result while in Section 3 we derive exponential bounds for the sufficient statistic. Proofs of our main statement and oracle inequality are given in Section 4. In Section 5 we explicit the constants and assumptions used in a list of classical law from the exponential family such as Poisson, exponential, gamma, beta, etc, and we study the particular case of categorical variables in Section 6 to assess the precision lost in dealing with general exponential family instead of directly bounding the particular distribution. Finally, in Section 7, we illustrate the performance of our approach on a simulation study based on the exponential distribution, and propose an application to DNA sequence distribution on a real data-set which we model using a piece-wise constant categorical distribution.

2. Model Selection Procedure

2.1. **Penalized maximum-likelihood estimator.** In our change-point setting, we will want to consider partitions m of the set $\{1, \dots, n\}$ on which our models will be piece-wise constant. More precisely, for a given partition m and denoting J a generic segment of m, we define the collection of models associated to m as:

$$\mathcal{S}_m = \{s_m | \forall J \in m, \forall t \in J, s_m(t) = \mathcal{G}(\boldsymbol{\theta}_J)\}.$$

We will consider the log-likelihood contrast γ and the associated Kullback-Leibler distance $K(s,u) = \mathbb{E}[\gamma(u) - \gamma(s)]$ between distribution s and u so that for distributions $s(t) = \mathcal{G}(\boldsymbol{\theta}_t) = h(Y_t)\exp\left[\boldsymbol{\theta}_t.\mathbf{T}_t - A(\boldsymbol{\theta}_t)\right]$, and $u = (\mathcal{G}(\boldsymbol{p}_t))_t$, we have

$$\begin{aligned} \gamma(s) &= \sum_{t=1}^{n} \left(A(\boldsymbol{\theta}_t) - \boldsymbol{\theta}_t.\mathbf{T}_t\right), \quad \text{and} \\ K(s,u) &= \sum_{t=1}^{n} \left[\nabla A(\boldsymbol{\theta}_t).(\boldsymbol{\theta}_t - \boldsymbol{p}_t) - (A(\boldsymbol{\theta}_t) - A(\boldsymbol{p}_t))\right]. \end{aligned}$$

The minimal contrast estimator $\hat{s}_m$ of s on the collection $\mathcal{S}_m$ is therefore $\hat{s}_m = \arg\min_{u\in\mathcal{S}_m} \gamma(u)$ and is given by

$$\hat{s}_m(t) \quad = \quad h(Y_t)\exp\left[\nabla A^{-1}(\bar{\mathbf{T}}_J).\mathbf{T}_t - A\left[\nabla A^{-1}(\bar{\mathbf{T}}_J)\right]\right],$$

where $\mathbf{T}_J = \sum_{t\in J}\mathbf{T}_t$, is the sum of the sufficient statistics on segment J, and $\bar{\mathbf{T}}_J = \mathbf{T}_J/|J|$ its mean; the bijective mapping of the gradient of A ensuring the existence and uniqueness of $\nabla A^{-1}(\bar{\mathbf{T}}_J)$.
Similarly, the projection $\bar{s}_m$ of s in terms of Kullback-Leibler on collection $\mathcal{S}_m$ is $\bar{s}_m = \arg\min_{u\in\mathcal{S}_m} K(s,u)$ and is given by

$$\bar{s}_m(t) \quad = \quad h(Y_t)\exp\left[\nabla A^{-1}(\bar{\mathbf{E}}_J).\mathbf{T}_t - A\left[\nabla A^{-1}(\bar{\mathbf{E}}_J)\right]\right],$$

where $\mathbf{E}_t = \nabla A(\boldsymbol{\theta}_t)$, $\mathbf{E}_J = \sum_{t\in J}\mathbf{E}_t$ and $\bar{\mathbf{E}}_J = \mathbf{E}_J/|J|$.

As is classical in penalized-likelihood settings, since minimizing the Kullback-Leibler risk would require knowing the true distribution s, we will wish to choose a penalty function $pen(m)$ such that the penalized estimator $\hat{s}_{\hat{m}}$, where $\hat{m} = \arg\min \gamma(\hat{s}_m) + pen(m)$, satisfies an oracle inequality of the type

$$\mathbb{E}[K(s,\hat{s}_{\hat{m}})] \leq C_1\mathbb{E}[K(s,\hat{s}_{m(s)})] + C_2$$

where C_2 is negligible compared to $C_1\mathbb{E}[K(s,\hat{s}_{m(s)})]$.

To this effect, here as in previous works (see for example[12]), we will introduce an event of large probability, $\Omega_{m_f}(\varepsilon)$ defined on a minimal partition m_f, where the fluctuation of each centered marginal is bounded. On this event we will derive tight controls of the risk of the models which will lead to the first part of the oracle inequality, *i.e.* $\mathbb{E}[K(s,\hat{s}_{\hat{m}})] \leq C_1\mathbb{E}[K(s,\hat{s}_{m(s)})]$. On the complementary of this event, we will obtain less tight results which are compensated on expectation by the negligible probability of the event. This control will result in the C_2 constant.
The choice of ε is therefore crucial in insuring that both C_1 and C_2 are as small as possible, while having the negligibility property of C_2 compared to $C_1\mathbb{E}[K(s,\hat{s}_{m(s)})]$. In practice, this choice is data-driven and is efficiently performed through the use of the slope heuristic [13]. In this paper, we therefore consider a generic but fixed ε, and aim at obtaining the shape of the penalty function. This choice is given in the next section.

2.2. **Main result.**

Theorem 2.1. *Let $s = \{s(t)\}_{1\leq t\leq n}$ be a distribution from the exponential family such that its parameters $\boldsymbol{\theta}_t$ belong to a compact set $\boldsymbol{\nu}$ with greatest band ν and such that $\nabla A(\boldsymbol{\nu}) \subset \{\mathbb{R}^+\}^d$. Assume there exist some positive constants ζ^i and κ^i and a positive constant κ such that*

$$\forall 1\leq t\leq n,\ \forall 1\leq i\leq d, \quad \zeta^i E_t^i \leq VarT_t^i \leq \kappa^i E_t^i \qquad (\mathbf{H})$$
$$\text{and} \quad \kappa \geq \max_t \max_i \{\kappa^i, 1/R_t^i, \alpha_t^i\sqrt{VarT_t^i}\}$$

where R_t^i is the radius of convergence of $A(\boldsymbol{\theta}_t + z\boldsymbol{\delta}_i)$, $\boldsymbol{\delta}_i = (0,\ldots,0,1,0,\ldots,0)$ and α_t^i is an exponential control of the growth rate of the coefficients of its power series.
Let $\zeta = min_i\{\zeta_i\}$. Let $\mathcal{M}_n$ be a collection of partitions constructed on a partition m_f such that there exists $\Gamma > 0$ satisfying $\forall J \in m_f, |J| \geq \Gamma\log(n)^2$, and let $(L_m)_{m\in\mathcal{M}_n}$ be some family of positive weights satisfying

$$\Sigma = \sum_{m\in\mathcal{M}_n} \exp(-L_m|m|) < +\infty.$$

Let $\varepsilon > 0$ and let β_ε be a positive constant depending on ε, distribution $\mathcal{G}$ and the observation y_t. If for every $m \in \mathcal{M}_n$

$$pen(m) \geq \beta_\varepsilon d|m|\left(1 + 4\sqrt{L_m}\right)^2,$$

then

$$\begin{aligned}\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] &\leq C_{\beta_\varepsilon}\inf_{m\in\mathcal{M}_n}\{K(s,\bar{s}_m)+pen(m)\}\\ &\quad +C(d,\nu,\kappa,\Gamma,\beta_\varepsilon,\Sigma,E^{max},E^{min}),\end{aligned}\tag{1}$$

where $E^{min}=\min_{1\leq i\leq d}\{\mathbb{E}(T^i)\}$ *and* $E^{max}=\max_{1\leq i\leq d}\{\mathbb{E}(T^i)\}$ *are bounds on the expectation of the sufficient statistics.*

Remark 2.2. *It can be noted that some assumptions of theorem 2.1 are in fact redundant and were included to simplify notations and clarify the dependency to the required assumptions. Indeed,*

- *the existence of* E^{min} *and* E^{max} *are guaranteed since we assume that* $\boldsymbol{\theta}$ *belongs to a compact* $\boldsymbol{\nu}$*. Similarly, the greatest bound* ν *of the compact needs not be introduced and we could use* $\theta^{max}-\theta^{min}$ *instead. However in general the marginals of the sufficient statistic might not fluctuate in the same subset of* $\mathbb{R}^+$ *and therefore* ν *is a refinement of the bound.*
- *bounding* $\boldsymbol{\theta}$ *(on* $\boldsymbol{\nu}$*) also implies that the variance of the sufficient statistics are bounded (since* A *is analytic), and therefore the existence of* ζ *and* κ *in assumption* (**H**) *are guaranteed.*

The theorem could therefore be re-written with the minimal conditions that $\boldsymbol{\theta}\in\boldsymbol{\nu}$, $\nabla A(\boldsymbol{\nu})\subset\{\mathbb{R}^+\}^d$ *and the hypothesis on* m_f *and* Σ*.*

In our change-point setting, we will choose weights L_m that depend on m only through their number of segments $|m|$, resulting in $L_m=1.1+\log\left(\frac{n}{|m|}\right)$ (see [1] for a justification of this choice). This leads to a penalty function of the form:

$$pen(m)=\beta_\varepsilon d|m|\left(1+4\sqrt{1.1+\log\left(\frac{n}{|m|}\right)}\right),\tag{2}$$

and the inequality of theorem 2.1 therefore becomes:

$$\begin{aligned}\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] &\leq C_{\beta_\varepsilon}\inf_{m\in\mathcal{M}_n}\left\{K(s,\bar{s}_m)+d\beta_\varepsilon|m|\left(1+4\sqrt{1.1+\log\left(\frac{n}{|m|}\right)}\right)\right\}\\ &+ C(d,\nu,\kappa,\Gamma,\beta_\varepsilon,\Sigma,E^{max},E^{min}).\end{aligned}$$

The following proposition gives a bound on the Kullback-Leibler risk associated to $\hat{s}_m$, which we prove in Section 4.1. This bound is then integrated into the main theorem to obtain our oracle inequality in corollary 2.4.

Proposition 2.3. *Under the assumptions of theorem 2.1, we have*

$$K(s,\bar{s}_m)+C(\varepsilon)d|m|-\frac{C(a,\Gamma,E^{min},E^{max},\varepsilon,\kappa)}{n^{a/2-1}}\leq\mathbb{E}[K(s,\hat{s}_m)],\tag{3}$$

where $a\geq 2$ *and* $C(\varepsilon)$ *is a constant that depends on* ε*, the constants involved in hypothesis* (**H**)*, and the constraints on the compact* $\boldsymbol{\nu}$*.*

Corollary 2.4. *Let* $s=\{s(t)\}_{1\leq t\leq n}$ *be a distribution from the exponential family such that its parameters* $\boldsymbol{\theta}_t$ *belong to a compact* $\boldsymbol{\nu}$ *of* $\{\mathbb{R}^+\}^d$*. Let* $\mathcal{M}_n$ *be a collection of partitions constructed on a partition* m_f *such that there exists* $\Gamma>0$ *satisfying* $\forall J\in m_f, |J|\geq\Gamma\log(n)^2$*, and assume* $\exists\ \zeta$ *and* κ *such that assumption* (**H**) *is verified. There exists some constant* C *such that*

$$\begin{aligned}\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] &\leq C\log(n)\inf_{m\in\mathcal{M}_n}\{\mathbf{E}[K(s,\hat{s}_m)]\}\\ &\quad +C(d,\kappa,\nu,\Gamma,E^{min},E^{max},\beta_\varepsilon,\Sigma).\end{aligned}\tag{4}$$

3. Exponential bounds

Following previous works [14, 15, 2], we will need to obtain exponential bounds on the fluctuation of the variables in order to derive bounds on the risk of models. Starting from equation

$$\gamma(\hat{s}_{\hat{m}}) + pen(\hat{m}) \quad \leq \quad \gamma(\hat{s}_m) + pen(m) \leq \gamma(\bar{s}_m) + pen(m),$$

and introducing the centered loss $\bar{\gamma}(u) = \gamma(u) - \mathbb{E}[\gamma(u)]$, we write

$$K(s, \hat{s}_{\hat{m}}) \quad \leq \quad K(s, \bar{s}_m) + \bar{\gamma}(\bar{s}_m) - \bar{\gamma}(\hat{s}_{\hat{m}}) - pen(\hat{m}) + pen(m).$$

We subsequently decompose $\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(\hat{s}_{m'})$ for any $m' \in \mathcal{M}_n$ in

$$\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(\hat{s}_{m'}) = (\bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'})) + (\bar{\gamma}(s) - \bar{\gamma}(\bar{s}_{m'})) + (\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)), \tag{5}$$

to control each term separately, with the main term now defined on the same model. This will lead us to introduce the chi-square statistic

$$\chi^2_m \;=\; \sum_{i=1}^{d} \chi^2(m,i) \;=\; \sum_{i=1}^{d} \sum_{J \in m} \frac{(T_J^i - E_J^i)^2}{E_J^i}, \tag{6}$$

and control the fluctuations of the sufficient statistics from their expectation.

3.1. Control of T_J.

We are now interested in controlling the variables T_J^i so that we can later apply Bernstein's inequality to $\chi^2(m,i)$. To this purpose, we wish to apply the large deviation result from Barraud and Birgé (see lemma 3 in [16]) which requires the control of the Laplace transform of T_J^i.

Let $i \in \{1, \ldots, d\}$ and $\boldsymbol{\delta}_i$ be the vector of $\mathbb{R}^d$ with 1 as the ith component and 0 as others. Let $1 \leq t \leq n$ and let z belong to $[-R_t^i, R_t^i]$, where R_t^i is the radius of convergence of the power series of $A(\boldsymbol{\theta}_t + z\boldsymbol{\delta}_i)$. We have

$$\log \mathbb{E}\left[e^{z(T_t^i - E_t^i)}\right] \quad = \quad \sum_{k \geq 2} c_k \frac{z^k}{k!} \;=\; c_2 \frac{z^2}{2} \sum_{k \geq 0} 2 \frac{c_{k+2}}{c_2} \frac{z^k}{(k+2)!},$$

where c_k is the kth cumulant of random variable T_t^i. By analyticity of A on the natural parameter space (see Lemma 3.3 [10] and [8]), there exists $\alpha_t^i > 0$ $\forall 1 \leq i \leq d$ such that

$$\forall k \geq 3, \; \left| \frac{c_k(T_t^i)}{c_2 (T_t^i)^{k/2}} \right| \leq \frac{1}{2} k! (\alpha_t^i)^{k-2},$$

and therefore if $\tau_i = \max_t \left\{ \max\{\alpha_t^i \sqrt{Var(T_t^i)}, 1/R_t^i\} \right\}$, the Laplace transform of T_t^i can be bounded as

$$\log \mathbb{E}\left[e^{z(T_t^i - E_t^i)}\right] \leq Var(T_t^i) \frac{z^2}{2} \sum_{k \geq 0} (\tau_i z)^k \leq \begin{cases} \dfrac{z^2}{2} \dfrac{\kappa^i E_t^i}{1 - \tau_i z} & \text{if } 0 \leq z < 1/\tau_i \\ \dfrac{z^2}{2} \kappa^i E_t^i & \text{if } -1/\tau_i < z < 0 \end{cases}.$$

We can therefore apply the large deviation result from [16]:

$$P\left[\sum_{t \in J} (T_t^i - E_t^i) \geq \sqrt{2\kappa^i x E_J^i} + \tau_i x \right] \leq e^{-x}$$

$$P\left[-\sum_{t \in J} (T_t^i - E_t^i) \leq \sqrt{2\kappa^i x E_J^i} \right] \leq e^{-x}$$

which implies, denoting $\kappa = \max_i \max\{\kappa^i, \tau_i\}$, that

$$P\left[\left| (T_J^i - E_J^i) \right| \geq x \right] \leq 2e^{-\frac{x^2}{2\kappa(E_J^i + x)}}.$$

3.2. **Exponential bound for χ^2_m.** We now consider m_f a partition of $\{1,\dots,n\}$ such that $\forall J \in m_f, |J| \geq \Gamma(\log(n))^2$, and we assume that $\mathcal{M}_n$ is a set of partitions that are constructed on this grid m_f. We first introduce the following set Ω_{m_f} defined by:

$$\Omega_{m_f}(\varepsilon) \quad = \quad \bigcap_{1\leq i\leq d} \Omega^i_{m_f}(\varepsilon) = \bigcap_{1\leq i\leq d} \bigcap_{J\in m_f} \left\{ \left| \frac{T^i_J}{E^i_J} - 1 \right| \leq \varepsilon \right\}.$$

Using the previous control, we have

$$\begin{aligned} \mathbf{P}(\Omega_{m_f}(\varepsilon)^C) &\leq \sum_{i=1}^{d} \sum_{J\in m_f} P\left[\left| T^i_J - E^i_J \right| \geq \varepsilon E^i_J \right] \leq \sum_{i=1}^{d} \sum_{J\in m_f} 2e^{-\frac{\varepsilon^2 E^i_J}{2\kappa(1+\varepsilon)}} \\ &\leq \frac{C(d,\varepsilon,\Gamma,\kappa,E^{min},a)}{n^a}, \end{aligned}$$

with $a > 2$.

The following proposition gives an exponential bound for χ^2_m on the restricted event $\Omega_{m_f}(\varepsilon)$.

Proposition 3.1. *Let $Y_1,\dots,Y_n$ be independent random variables with distribution s (from the exponential family) and verifying (**H**). Let m be a partition of $\mathcal{M}_n$ with $|m|$ segments and χ^2_m be the statistic given by (6). For any positive x, we have*

$$\mathbf{P}\left[\chi^2_m \mathbf{1}_{\Omega_{m_f}(\varepsilon)} \geq d\ \kappa \left(|m| + 8(1+\varepsilon)\sqrt{x|m|} + 4(1+\varepsilon)x \right) \right] \quad \leq \quad d\ e^{-x}.$$

Proof: We have

$$\mathbb{E}[\chi^2(m,i)] = \sum_{J\in m} \frac{1}{E^i_J} \mathbb{E}(T^i_J - E^i_J)^2,$$

and therefore since $\zeta^i E^i_J \leq Var(T^i_J) \leq \kappa^i E^i_J$

$$\zeta|m| \leq \mathbb{E}[\chi^2(m,i)] \leq \kappa|m|.$$

We introduce the variables $Z_J(i)$ such that

$$\chi^2(m,i) = \sum_{J\in m} Z_J(i) = \sum_{J\in m} \frac{(T^i_J - E^i_J)^2}{E^i_J},$$

and control their moments using

$$\begin{aligned} \mathbf{E}\left[Z_J(i)^p \mathbf{1}_{\Omega_{m_f}(\varepsilon)} \right] &\leq \left(\frac{1}{E^i_J} \right)^p \int_0^{+\infty} x^p\ dP\left[\left\{ (T^i_J - E^i_J)^2 \geq x \right\} \cap \Omega_{m_f}(\varepsilon) \right] dx \\ &\leq \left(\frac{1}{E^i_J} \right)^p \int_0^{\varepsilon E^i_J} 2p\ x^{2p-1} P\left[|T^i_J - E^i_J| \geq x \right] dx \\ &\leq \left(\frac{1}{E^i_J} \right)^p \int_0^{\varepsilon E^i_J} 4p\ x^{2p-1} e^{-\frac{x^2}{2\kappa E^i_J(1+\varepsilon)}} dx \\ &\leq 4p\kappa^p\,(1+\varepsilon)^p \int_0^{+\infty} u^{2p-1} e^{-\frac{u^2}{2}} du \\ &\leq 4p\kappa^p\,(1+\varepsilon)^p \int_0^{+\infty} (2t)^{p-1}\, e^{-t} dt \\ &\leq 2^{p+1} p\kappa^p\,(1+\varepsilon)^p\, p! \end{aligned}$$

We can then conclude by applying Bernstein's inequality [12] taking $v = 2^5\left(\kappa(1+\varepsilon)\right)^2|m|$ and $c = 4\left(\kappa(1+\varepsilon)\right)$:

$$\mathbf{P}\left[\chi^2(m,i) \mathbf{1}_{\Omega_{m_f}(\varepsilon)} \geq \kappa \left(|m| + 8(1+\varepsilon)\sqrt{x|m|} + 4(1+\varepsilon)x \right) \right] \quad \leq \quad e^{-x}.$$

While for a given i the $\{Z_J(i)\}_{J\in m}$ are independent variables, in general for a given J the variables $\{Z_J(i)\}_{1\leq i\leq d}$ are not. We conclude the proof using lemma 3.2:

Lemma 3.2. *Let $X_1, \dots, X_n$ be real random variables and let $a_1, \dots, a_n \in \mathbb{R}^n$ and $b_1, \dots, b_n \in [0,1]^n$ such that $\forall 1 \leq i \leq n,\ \mathbb{P}(X_i \geq a_i) \leq b_i$. Then*

$$\mathbb{P}\left(\sum_{i=1}^{n} X_i \geq \sum_{i=1}^{n} a_i\right) \leq \sum_{i=1}^{n} b_i.$$

4. Proof of our main results

4.1. **Proof of proposition 2.3.** Before we focus on the main theorem, we prove our lower bound on the risk of a model. By definition, we have

$$K(\bar{s}_m, \hat{s}_m) = \sum_{J \in m} |J| \left[A\left(\nabla A^{-1}(\bar{\mathbf{T}}_J)\right) - A\left(\nabla A^{-1}(\bar{\mathbf{E}}_J)\right) - \bar{\mathbf{E}}_J . \left(\nabla A^{-1}(\bar{\mathbf{T}}_J) - \nabla A^{-1}(\bar{\mathbf{E}}_J)\right)\right].$$

We define the compact subset $K(\varepsilon)$ of $\mathbb{R}^d$ as the pre-image by ∇A of the domains induced by Ω_m, that is

$$K(\varepsilon) = \left\{ \mathbf{z} \in \mathbb{R}^d \,\middle|\, \nabla A(\mathbf{z}) \in \bigcup_{m \in \mathcal{M}} \bigcup_{J \in m} \mathcal{B}(\bar{E}_J, \varepsilon \bar{E}_J) \right\},$$

where $\mathcal{B}(\mathbf{x}, r)$ denotes the closed ball centered in $\mathbf{x}$ with radius r of $\mathbb{R}^d$. Since we consider the union of a finite number of balls, homeomorphic properties of ∇A ensures that $K(\varepsilon)$ is a compact set of $\mathbb{R}^d$.
Since $\nabla^2 A$ is definite positive, A is m_ε-strongly convex on the compact set $K(\varepsilon)$, and we have

$$\begin{aligned} K(\bar{s}_m, \hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)} &\geq \sum_{J \in m} |J| \frac{m_\varepsilon}{2} ||\nabla A^{-1}(\bar{\mathbf{T}}_J) - \nabla A^{-1}(\bar{\mathbf{E}}_J)||^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)} \\ &\geq \sum_{i=1}^{d} \sum_{J \in m} |J| \frac{m_\varepsilon}{2E_J^i} E_J^i \left[\nabla A^{-1}(\bar{\mathbf{T}}_J)^i - \nabla A^{-1}(\bar{\mathbf{E}}_J)^i\right]^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)}, \end{aligned}$$

where m_ε can be chosen as a lower bound on the smallest eigen-value of $\nabla^2 A$ on $K(\varepsilon)$.
Now introducing

$$V_m^2 = \sum_{i=1}^{d} V^2(m,i) = \sum_{i=1}^{d} \sum_{J \in m} E_J^i \left[\nabla A^{-1}(\bar{\mathbf{E}}_J)^i - \nabla A^{-1}(\bar{\mathbf{T}}_J)^i\right]^2, \tag{7}$$

we obtain

$$K(\bar{s}_m, \hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)} \geq \frac{m_\varepsilon}{2E^{max}} V_m^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)}. \tag{8}$$

Now if we consider M_ε an upper bound of the eigen-values of $\nabla^2 A$ on $K(\varepsilon)$, then on $\Omega_{m_f}(\varepsilon)$ we have

$$||\nabla A^{-1}(\bar{\mathbf{T}}_J) - \nabla A^{-1}(\bar{\mathbf{E}}_J)||^2 \geq \frac{1}{M_\varepsilon} ||\bar{\mathbf{T}}_J - \bar{\mathbf{E}}_J||^2,$$

and therefore

$$\begin{aligned} K(\bar{s}_m, \hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)} &\geq \sum_{J \in m} |J| \frac{m_\varepsilon}{2M_\varepsilon} ||\bar{\mathbf{T}}_J - \bar{\mathbf{E}}_J||^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)} \\ &\geq \frac{E^{min} m_\varepsilon}{2M_\varepsilon} \chi_m^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)}. \end{aligned}$$

Combining the previous results leads to

$$\frac{1}{2} E^{min} \frac{m_\varepsilon}{M_\varepsilon}\ \mathbb{E}[\chi_m^2] - \mathbb{E}[\chi_m^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)^c}] \leq \mathbb{E}[K(\bar{s}_m, \hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)}],$$

and using $\mathbb{E}[\chi_m^2] \geq d\zeta|m|$ and Cauchy-Schwarz inequality on $\mathbb{E}[\chi_m^2 \mathbf{1}_{\Omega_{m_f}(\varepsilon)^c}]$,

$$\frac{1}{2} \xi E^{min} \frac{m_\varepsilon}{M_\varepsilon}\ d|m| - \frac{C(a, \Gamma, E^{min}, E^{max}, \varepsilon, \kappa, d)}{n^{a/2-1}} \leq \mathbb{E}[K(\bar{s}_m, \hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)}].$$

Finally, introducing $C(\varepsilon) = \frac{1}{2}\xi E^{min}\frac{m_\varepsilon}{M_\varepsilon}$, and since $\mathbb{E}\left[K(\bar{s}_m, \hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)^C}\right] \geq 0$, we have

$$K(s, \bar{s}_m) + C(\varepsilon)\ d|m| - \frac{C(a, \Gamma, E^{min}, E^{max}, \varepsilon, \kappa, d)}{n^{a/2-1}} \leq \mathbb{E}[K(s, \hat{s}_m)].$$

4.2. **Proof of theorem 2.1.** To prove theorem 2.1, we will control each of the three terms in equation (5) individually. To this effect, we introduce a generic partition m' of $\mathcal{M}_n$.

4.2.1. *Control of term* $\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})$. Let

$$\Omega_1(\xi) = \bigcap_{m' \in \mathcal{M}_n} \Big\{ \chi^2_{m'} \mathbf{1}_{\Omega_{m_f}(\varepsilon)} \leq d\kappa \,[|m'| \quad + \quad 8(1+\varepsilon)\sqrt{(L_{m'}|m'| + \xi)|m'|} \\ + \quad 4(1+\varepsilon)(L_{m'}|m'| + \xi)]\Big\}.$$

Then we can control $\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})$ with the following proposition :

Proposition 4.1.

$$\left(\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})\right) \mathbf{1}_{\Omega_{m_f}(\varepsilon) \cap \Omega_1(\xi)} \leq \frac{1}{1+\varepsilon} K(\bar{s}_{m'}, \hat{s}_{m'}) + \\ + \frac{dE^{max}\kappa(1+\varepsilon)}{2m_\varepsilon}\left[|m'| + 8(1+\varepsilon)\sqrt{(L_{m'}|m'| + \xi)|m'|} + 4(1+\varepsilon)(L_{m'}|m'| + \xi)\right]$$

Proof: We recall that

$$\begin{aligned} \bar{\gamma}(\bar{s}_{m'}) &= \sum_{J \in m'} |J| \left[-\nabla A^{-1}(\bar{\mathbf{E}}_J).(\bar{\mathbf{T}}_J - \bar{\mathbf{E}}_J)\right] \text{ and} \\ \bar{\gamma}(\hat{s}_{m'}) &= \sum_{J \in m'} |J| \left[-\nabla A^{-1}(\bar{\mathbf{T}}_J).(\bar{\mathbf{T}}_J - \bar{\mathbf{E}}_J)\right]. \end{aligned}$$

We therefore have

$$\begin{aligned} \bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'}) &= \sum_{J \in m'} |J| \left[(\bar{\mathbf{E}}_J - \bar{\mathbf{T}}_J).\left(\nabla A^{-1}(\bar{\mathbf{T}}_J) - \nabla A^{-1}(\bar{\mathbf{E}}_J)\right)\right] \\ &= \sum_{i=1}^{d} \sum_{J \in m'} |J| \left[(\bar{E^i}_J - \bar{T^i}_J)\left(\nabla A^{-1}(\bar{\mathbf{T}}_J)^i - \nabla A^{-1}(\bar{\mathbf{E}}_J)^i\right)\right]. \end{aligned}$$

Using Cauchy-Schwarz inequality,

$$\begin{aligned} \bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'}) &= \sum_{i=1}^{d} \sum_{J \in m'} \left[\frac{(E^i{}_J - T^i{}_J)}{\sqrt{E^i_J}} \sqrt{E^i_J}\left(\nabla A^{-1}(\bar{\mathbf{T}}_J)^i - \nabla A^{-1}(\bar{\mathbf{E}}_J)^i\right)\right] \\ &\leq \left[\sqrt{\sum_{i=1}^{d} \sum_{J \in m'} \frac{(T^i{}_J - E^i{}_J)^2}{E^i_J}} \sqrt{\sum_{i=1}^{d} \sum_{J \in m'} E^i_J \left[\nabla A^{-1}(\bar{\mathbf{T}}_J)^i - \nabla A^{-1}(\bar{\mathbf{E}}_J)^i\right]^2}\right] \\ &\leq \sqrt{\chi^2_{m'}}\sqrt{V^2_{m'}}, \end{aligned}$$

where χ^2_m has been defined in equation (6) and V^2_m has been defined in equation (7). Then from equation (8) and $2ab \leq xa^2 + x^{-1}b^2$, we get on $\Omega_{m_f}(\varepsilon)$

$$\bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'}) \leq \frac{1+\varepsilon}{2m_\varepsilon} E^{max} \chi^2_{m'} + \frac{1}{1+\varepsilon} K(\bar{s}_{m'}, \hat{s}_{m'}).$$

Finally, using proposition 3.1,

$$\begin{aligned} \left(\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})\right) \mathbf{1}_{\Omega_{m_f}(\varepsilon) \cap \Omega_1(\xi)} &\leq \frac{d\kappa(1+\varepsilon)}{2m_\varepsilon} E^{max} \Big[|m'| + 8(1+\varepsilon)\sqrt{(L_{m'}|m'| + \xi)|m'|} \\ &\quad + 4(1+\varepsilon)(L_{m'}|m'| + \xi)] + \frac{1}{1+\varepsilon} K(\bar{s}_{m'}, \hat{s}_{m'}). \end{aligned}$$

4.2.2. *Control of term* $\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)$. The expectation of the second term can be bounded using the following proposition:

Proposition 4.2.

$$\left|\mathbb{E}\left[(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s))\mathbf{1}_{\Omega_{m_f}(\varepsilon)}\right]\right| \leq \frac{C(d,\varepsilon,\Gamma,\kappa,\nu,E^{min},E^{\max})}{n^{(a-1)/2}},$$

where ν *is the greatest length of the compact* $\boldsymbol{\nu}$.

Proof: We recall that

$$\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s) = \sum_{J\in m}\sum_{t\in J}(\boldsymbol{\theta}_t - \nabla A^{-1}(\bar{\mathbf{E}}_J)).(\mathbf{T}_t - \mathbf{E}_t).$$

Then

$$\begin{aligned}
\left|\mathbb{E}\left[(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s))\mathbf{1}_{\Omega_{m_f}(\varepsilon)}\right]\right| &\leq \mathbb{E}\left[|\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)|\,\mathbf{1}^C_{\Omega_{m_f}(\varepsilon)}\right] \\
&\leq \mathbb{E}\left[\left|\sum_{J\in m}\sum_{t\in J}\sum_{i=1}^{d}(\theta_t^i - \nabla A^{-1}(\bar{\mathbf{E}}_J)^i)(T_t^i - E_t^i)\right|\mathbf{1}_{\Omega_{m_f}(\varepsilon)^C}\right] \\
&\leq \nu\sum_{i=1}^{d}\left[\mathbb{E}\left(\sum_{t=1}^{n}(T_t^i - E_t^i)\right)^2\right]^{1/2}\mathbb{P}\left(\Omega^i_{m_f}(\varepsilon)^C\right)^{1/2} \\
&\leq \nu\sum_{i=1}^{d}\left[\sum_{t=1}^{n}\kappa E_t^i\right]^{1/2}\mathbb{P}\left(\Omega^i_{m_f}(\varepsilon)^C\right)^{1/2} \\
&\leq d\nu\frac{\sqrt{2\kappa E^{\max}C(\varepsilon,\Gamma,\kappa,E^{min})}}{n^{(a-1)/2}}.
\end{aligned}$$

4.2.3. *Control of term* $\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})$. Finally, we control the last term using proposition 4.3 for which a proof can be found in [2].

Proposition 4.3.

$$\mathbf{P}\left[\bar{\gamma}(s) - \bar{\gamma}(u) \geq K(s,u) - 2h^2(s,u) + 2dx\right] \leq e^{-dx} \leq de^{-x}.$$

Applying it to $u = \bar{s}_{m'}$ yields:

$$\mathbf{P}\left[\bar{\gamma}(s) - \bar{\gamma}(\bar{s}_{m'}) \geq K(s,\bar{s}_{m'}) - 2h^2(s,\bar{s}_{m'}) + 2dx\right] \leq de^{-x}.$$

We then define

$$\Omega_2(\xi) = \bigcap_{m'\in\mathcal{M}_n}\left\{\bar{\gamma}(s) - \bar{\gamma}(\bar{s}_{m'}) \leq K(s,\bar{s}_{m'}) - 2h^2(s,\bar{s}_{m'}) + 2d(L_{m'}|m'| + \xi)\right\}.$$

4.2.4. *Proof of the theorem.* We can now combine the previous propositions in order to prove our main theorem. To this effect, we introduce the following event: $\Omega(\varepsilon,\xi) = \Omega_{m_f}(\varepsilon)\cap\Omega_1(\xi)\cap\Omega_2(\xi)$. On this event, we have, with $R = \bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)$:

$$\begin{aligned}
(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(\hat{s}_{\hat{m}}))\mathbf{1}_{\Omega(\varepsilon,\xi)} &= (\bar{\gamma}(s) - \bar{\gamma}(\bar{s}_{\hat{m}}))\mathbf{1}_{\Omega(\varepsilon,\xi)} + (\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s))\mathbf{1}_{\Omega(\varepsilon,\xi)} + (\bar{\gamma}(\bar{s}_{\hat{m}}) - \bar{\gamma}(\hat{s}_{\hat{m}}))\mathbf{1}_{\Omega(\varepsilon,\xi)} \\
&\leq \left[K(s,\bar{s}_{\hat{m}}) - 2h^2(s,\bar{s}_{\hat{m}})\right]\mathbf{1}_{\Omega(\varepsilon,\xi)} + R\mathbf{1}_{\Omega(\varepsilon,\xi)} + \frac{1}{1+\varepsilon}K(\bar{s}_{\hat{m}},\hat{s}_{\hat{m}})\mathbf{1}_{\Omega(\varepsilon,\xi)} \\
&\quad + \frac{dE^{max}\kappa(1+\varepsilon)}{2m_\varepsilon}\left[|\hat{m}| + 8(1+\varepsilon)\sqrt{(L_{\hat{m}}|\hat{m}| + \xi)|\hat{m}|} + 4(1+\varepsilon)(L_{\hat{m}}|\hat{m}| + \xi)\right] \\
&\quad + 2dL_{\hat{m}}|\hat{m}| + 2d\xi.
\end{aligned}$$

We then apply the proof of [2] with $C(\varepsilon) = \dfrac{E^{max}\kappa(1+\varepsilon)}{2m_\varepsilon}$, which yields:

$$\begin{aligned}
\frac{\varepsilon}{1+\varepsilon}h^2(s,\hat{s}_{\hat{m}})\mathbf{1}_{\Omega(\varepsilon,\xi)} &\leq K(s,\bar{s}_m)\mathbf{1}_{\Omega(\varepsilon,\xi)} + R\mathbf{1}_{\Omega(\varepsilon,\xi)} - pen(\hat{m}) + pen(m) \\
&\quad + d|\hat{m}|C_2(\varepsilon)\left(1 + 4\sqrt{L_{\hat{m}}}\right)^2 + 2\xi d\left[1 + (1+\varepsilon)C(\varepsilon)\left(\frac{8}{\varepsilon} + 2\right)\right],
\end{aligned}$$

where $C_2(\varepsilon) = \dfrac{E^{max}\kappa}{2m_\varepsilon}(1+\varepsilon)^3$. Then since by assumption, $pen(\hat{m}) \geq \beta_\varepsilon d|\hat{m}|\left(1+4\sqrt{L_{\hat{m}}}\right)^2$, choosing $\beta_\varepsilon = C_2(\varepsilon)$ yields:

$$h^2(s,\hat{s}_{\hat{m}})\mathbf{1}_{\Omega(\varepsilon,\xi)} \quad \leq \quad C_{\beta_\varepsilon}\left[K(s,\bar{s}_m)\mathbf{1}_{\Omega(\varepsilon,\xi)} + R\mathbf{1}_{\Omega(\varepsilon,\xi)} + pen(m)\right] + d\xi C(\beta_\varepsilon),$$

with $C_{\beta_\varepsilon} = \dfrac{1+C_2^{-1}(\beta_\varepsilon)}{C_2^{-1}(\beta_\varepsilon)}$.

From propositions 3.1 and 4.3 comes both $\mathbf{P}\left(\Omega_1(\xi)^C\right) \leq d\sum_{m'\in\mathcal{M}_n} e^{-(L_{m'}|m'|+\xi)}$ and $\mathbf{P}\left(\Omega_2(\xi)^C\right) \leq d\sum_{m'\in\mathcal{M}_n} e^{-(L_{m'}|m'|+\xi)}$, so that using hypothesis (1),

$$\mathbf{P}\left(\Omega_1(\xi)^C \cup \Omega_2(\xi)^C\right) \quad \leq \quad 2d\sum_{m'\in\mathcal{M}_n} e^{-(L_{m'}|m'|+\xi)} \leq 2d\Sigma e^{-\xi},$$

and thus $\mathbf{P}\left(\Omega_1(\xi)\cap\Omega_2(\xi)\right) \geq 1-2d\Sigma e^{-d\xi}$. Integrating over ξ and using proposition 4.2 leads to

$$\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\mathbf{1}_{\Omega_{m_f}(\varepsilon)}\right] \quad \leq \quad C_{\beta_\varepsilon}\left[K(s,\bar{s}_m) + \frac{C(d,\nu,\beta_\varepsilon,\Gamma,\kappa,E^{min},E^{\max},a)}{n^{(a-1)/2}} + pen(m)\right] + 2d\Sigma C(\beta_\varepsilon).$$

And since $\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\mathbf{1}_{\Omega_{m_f}(\varepsilon)^C}\right] \leq \dfrac{C(d,\beta_\varepsilon,\Gamma,\kappa,E^{min},a)}{n^{a-1}}$, we have

$$\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] \quad \leq \quad C_{\beta_\varepsilon}\left[K(s,\bar{s}_m) + pen(m)\right] + C'(d,\nu,\beta_\varepsilon,\Gamma,\kappa,E^{min},E^{\max},\Sigma).$$

Finally, by minimizing over $m\in\mathcal{M}_n$, we get

$$\mathbf{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] \quad \leq \quad C_{\beta_\varepsilon}\inf_{m\in\mathcal{M}_n}\left\{K(s,\bar{s}_m) + pen(m)\right\} + C'(d,\nu,\beta_\varepsilon,\Gamma,\kappa,E^{min},E^{\max},\Sigma).$$

5. Characteristics of classic laws

In this section we provide some explicit values for the constants used in the main theorem for some usual distributions from the exponential family. While we do not detail all computations, in Table 1 we summarize the parameters of the distribution, the sufficient statistics associated with the natural parameters, the expression of the log-partition function and possible choices for the values of ζ and κ.

Here we detail the computations for the exponential distribution. Some details for other distributions are provided in Appendix A.
The exponential distribution can be written in its natural form as

$$s(t) = \exp\left(-\lambda_t Y_t + \log\lambda_t\right)$$

from which we obtain, dropping the index t

$$\begin{aligned} \theta &= -\lambda \\ A(\theta) &= -\log(-\theta) \\ T(Y) &= Y. \end{aligned}$$

- Variance / Expectation relationship
  $A'(\theta) = -\frac{1}{\theta}$ and $A"(\theta) = \frac{1}{\theta^2}$ so that
  $$-\frac{1}{\theta^{min}}E \leq Var(T) \leq -\frac{1}{\theta^{max}}E.$$
- Cumulants
  $A(\theta+z) - A(\theta) = -\log\left(1+\frac{z}{\theta}\right)$ is analytic for $-|\theta| < z < |\theta|$, *i.e.* $R = |\theta|$. We then obtain
  $$A(\theta+z) - A(\theta) = \sum_{k\geq 1}(-1)^k\frac{z^k}{k\theta^k} = \sum_{k\geq 1}\left(-\frac{1}{\theta}\right)^k(k-1)!\frac{z^k}{k!}$$
  so that for $k\geq 1, c_k = \left(-\frac{1}{\theta}\right)^k(k-1)!$.

- Bounds
  We have, for $k \geq 3$;
$$\left|\frac{c_k}{c_2^{k/2}}\right| = (k-1)! \leq \frac{k}{2}(k-1)! \leq \frac{1}{2}k!(\alpha)^{k-2}$$
  with $\alpha = 1$. Then $\tau = \max\left\{\alpha\sqrt{Var(T)}, \frac{1}{R}\right\} = \max\left\{\frac{1}{|\theta|}, \frac{1}{|\theta|}\right\} = -\frac{1}{\theta^{max}}$ and finally, $\zeta = -\frac{1}{\theta^{min}}$ and $\kappa = -\frac{1}{\theta^{max}}$.

Distribution	Natural Parameters	Sufficient Statistics	Log-Partition	ζ	κ
Poisson	$\log\lambda$	x	e^θ	1	1
Exponential	$-\lambda$	x	$-\log(-\theta)$	$-\frac{1}{\theta^{min}}$	$-\frac{1}{\theta^{max}}$
Gaussian ($\mu > 0$)	$\frac{\mu}{\sigma^2}, -\frac{1}{2\sigma^2}$	(x, x^2)	$-\frac{{\theta_1}^2}{4\theta_2} - \frac{1}{2}\log(-2\theta_2)$	$\min\left\{\frac{1}{\theta_1^{max}}, -\frac{1}{\theta_2^{min}}\right\}$	$\max\left\{\frac{1}{\theta_1^{min}}, -\frac{2}{\theta_2^{max}}\right\}$
Pareto ($\alpha < \frac{-1}{log x_m}$)	$-\alpha$	$\log x$	$\theta\log x_m - \log(-\theta)$	$\frac{1}{\theta^{max}}\frac{1}{\theta^{max}\log x_m - 1}$	$\frac{\log x_m}{\theta^{min}\log x_m - 1}$
Gamma (α fixed)	$-\beta$	x	-$\alpha\log(-\theta)$	$-\frac{1}{\theta^{min}}$	$-\frac{1}{\theta^{max}}$
Weibull (k fixed)	$-\lambda^{-k}$	x^k	$-\log(-\theta)$	$-\frac{1}{\theta^{min}}$	$-\frac{1}{\theta^{max}}$
Laplace (μ fixed)	$-1/b$	$\lvert x-\mu\rvert$	$-\log(-\theta)$	$-\frac{1}{\theta^{min}}$	$-\frac{1}{\theta^{max}}$
Binomial (n fixed)	$\log(\frac{p}{1-p})$	x	$n\log(1+e^\theta)$	$\frac{1}{1+e^{\theta_{max}}}$	2
Negative Binomial (ϕ fixed)	$\log(1-p)$	x	$-\phi\log(1-e^\theta)$	1	$e^{-\theta_{min}}$

TABLE 1. List of most usual univariate laws from the exponential family.

6. Particular case of categorical variables

The goal of this section is double. We first wish to illustrate how our general approach can be used when working with a specific distribution, in this example the categorical distribution. Indeed, while we have shown that the penalty shape is the same for all distribution, the constants can be defined properly with quantities depending on the distribution of interest. We then show how the main results could be derived if working directly with the characteristics of the categorical distribution instead of dealing with the log-partition function from the exponential family. We conclude this section by comparing the constants obtained in each case.

6.1. **Application of the general approach.** Here we suppose that Y can take values between 1 and $d+1$ and we denote $\{p^i; 1 \leq i \leq d\}$ the probability that Y belongs to categories 1 through d (so that $p_{d+1} = 1 - \sum_{i=1}^d p^i$).

In the canonical form, the parameters $\boldsymbol{\theta}$ are given by $\theta^i = \log\frac{p^i}{1-\sum_{i=1}^d p^i}$ (for $1 \leq i \leq d$) and we have

- $T^i = \mathbf{1}_{\{Y=i\}}$
- $A(\boldsymbol{\theta}) = \log(1 + \sum_{i=1}^d e^{\theta^i})$

- $E^i = \dfrac{e^{\theta^i}}{1+\sum_{j=1}^d e^{\theta^j}}$ and

$$\nabla^2 A(\boldsymbol{\theta}) = \begin{pmatrix} E^1(1-E^1) & -E^1E^2 & \dots & -E^1E^d \\ -E^1E^2 & E^2(1-E^2) & \dots & -E^2E^d \\ \vdots & & & \\ -E^1E^d & \dots & & E^d(1-E^d) \end{pmatrix}$$

The inverse mapping of the gradient of the log-partition function is given by $[(\nabla A)^{-1}(\boldsymbol{\beta})]^i = \log \dfrac{\beta^i}{1-\sum_{j=1}^d \beta^j}$.

6.1.1. *Verification of hypothesis (H).* Let us assume that the probabilities of categories 1 to d are bounded away from 0 and 1. We will want to assume the same property on the *normalizing* category $d+1$, so that there exists $a > 1$ such that $-\log a < \theta_i < \log a$ for all $1 \le i \le d$. Our compact set $\boldsymbol{\nu}$ is therefore defined by $[-\log a; \log a]^d$, and $\nu = 2\log a$. This leads to:

$$E^{min} = \frac{1}{a+a^2 d} > 0 \qquad \text{and} \qquad E^{max} = \frac{a^2}{d+a} < 1.$$

Now since for all i we have $Var(T^i) = (1-E^i)E^i$, we get $\dfrac{d+a(1-a)}{d+a} E^i \le VarT^i \le E^i$ and therefore

$$\zeta = \frac{d+a(1-a)}{d+a}.$$

The Laplace transform of sufficient statistic T^i is

$$\log \mathbb{E}\left[e^{z(T^i-E^i)}\right] \quad = \quad \log\left[\frac{1+\sum_j e^{\theta^j} + e^{\theta^i}(e^z-1)}{1+\sum_j e^{\theta^j}}\right] = \log\left[1+E^i(e^z-1)\right],$$

which is analytic in z provided $z \le \log\left(1+\frac{1}{E^i}\right)$.

Let $R = \log 2$ and let $z \in [-R; R]$. Then the cumulants associated with T^i can be obtained through the recurrence property $c_1^i = p^i$ and for $k \ge 2$,

$$c_{k+1}^i = p^i(1-p^i)\frac{\partial c_k^i}{\partial p^i},$$

where we have switched back to usual proportion notations for sake of readability.
Let $P_{k+1}(p) = \frac{\partial c_k^i}{\partial p}$ with $p = p^i$. We have $P_2(p) = 1$, polynomial in p of degree $0 = k-2$ and with sum of absolute coefficients $= 1 = 2^{2-3} * 2!$. Let us assume that for a given $k > 1$, $P_k(p)$ is a polynomial in p of degree $k-2$ and with sum of absolute coefficients less than $2^{k-3}k!$. Then

$$P_{k+1}(p) = \frac{\partial c_k^i}{\partial p} = \frac{\partial\left[p(1-p)P_k(p)\right]}{\partial p} = (1-2p)P_k(p) + (p-p^2)\frac{\partial P_k(p)}{\partial p}$$

Thus denoting $P_k(p) = \sum_{\ell=0}^{k-2} a_\ell p^\ell$ we get $P_{k+1}(p) = a_0 + \sum_{\ell=1}^{k-2}(\ell+1)(a_\ell - a_{\ell-1})p^\ell - k a_{k-2}p^{k-1}$ which is of degree $k-1$ and with sum of absolute coefficient less than

$$\begin{aligned} |a_0| + \sum_{\ell=1}^{k-2}(\ell+1)|a_\ell - a_{\ell-1}| + k|a_{k-2}| \quad &\le \quad 3|a_0| + \sum_{\ell=1}^{k-3}(2*\ell+3)|a_\ell| + (2*k-1)|a_{k-2}| \\ &\le \quad \sum_{\ell=0}^{k-2}(2*\ell+3)|a_\ell| \le 2*(k+1)\sum_{\ell=0}^{k-2}|a_\ell| \le 2^{k-2}(k+1)! \end{aligned}$$

We therefore obtain by recurrence that P_k (for $k \ge 2$) is a polynomial in p of degree $k-2$ and with sum of absolute coefficients less than $2^{k-3}k!$. From this we get:

$$\left|\frac{c_k(T^i)}{c_2(T^i)^{k/2}}\right| \le \frac{1}{2}k!\left(\frac{2}{\sqrt{p^i(1-p^i)}}\right)^{k-2}$$

and $\tau_i = \max\left\{\frac{2}{\sqrt{p^i(1-p^i)}} * \sqrt{p^i(1-p^i)}, \frac{1}{R^i}\right\} = \max\{2, 1/R\}$. Since $\kappa_i = 1$ and $R = \log 2$, we can take $\kappa = 2$.

Note that considering the sufficient statistic T^i independently of other categories can be interpreted as considering a surrogate random variable X with success $X = 1$ if $Y = 1$ and $X = 0$ otherwise. The distribution of X is then Bernoulli with parameter $p = p^i$ and the results claimed in the previous Section can be deducted from what precedes.

6.1.2. *Computation of m_ε and M_ε.* For a given vector of proportions $(\pi_1, \ldots, \pi_d)$, the smallest eigen value of $\nabla^2 A$ is greater than $\Pi = (1 - \sum_{i=1}^d \pi_i) \prod_{i=1}^d \pi_i$. This can be seen by applying Sylvester's criterion to the symmetric matrix $\nabla^2 A - \Pi Id$. Considering the mapping between natural parameters and proportions, we obtain that Π can be defined as $\Pi(\eta) = \dfrac{\exp(\sum \eta^i)}{\left[1 + \sum e^{\eta^i}\right]^{d+1}}$.

Therefore let $z \in K(\varepsilon)$. There exists E such that $\nabla A(z) \in \mathcal{B}(E, \varepsilon)$, that is there exists E such that $\forall i, \frac{e^{z^i}}{1+\sum e^{j^j}} \in [E^i - \varepsilon; E^i + \varepsilon]$. We therefore have:

$$\Pi(z) = \frac{\prod e^{z^i}}{\left[1 + \sum e^{z^i}\right]^{d+1}} \geq (E^{min} - \varepsilon)^{d+1}$$

and can conclude by taking $m_\varepsilon = (E^{min} - \varepsilon)^{d+1}$.

The greatest eigen-value of $\nabla^2 A$ can be bounded by the trace of the gradient matrix, in this case the sum of the variance of each sufficient statistic. Therefore, for a distribution with proportions $\{q_i\}$, one would get

$$M \leq \sum_{i=1}^d q_i(1 - q_i) \leq \sum_{i=1}^d q_i \leq 1.$$

Therefore, on any set we have $M \leq 1$ and in particular, on $K(\varepsilon)$, we can take $M_\varepsilon = 1$.

6.2. **Direct results for the categorical variables.** This section is dedicated to the derivation of direct controls when studying the categorical distribution. As before, Y will take values in $1, \ldots, d+1$ and p^i will denote the probability of category i. Once again, it is possible to reduce the number of parameters to d, however keeping all $d+1$ parameters leads to more tractable quantities to control and to smaller resulting constants. For sake of readability and to avoid confusions, we will denote $r = d + 1$.

6.2.1. *Notations and main result.* The density s_t of Y_t can be decomposed in r terms by denoting $s(t, i) = P(Y_t = i) = p_t^i = E[\mathbf{1}_{\{Y=i\}}]$. In this specific case, quantities such as the contrast or the Kullback-Leibler risk can be identified :

- The model S_m is:
$$\mathcal{S}_m = \left\{ \begin{array}{c} u : \{1, ..., n\} \times \{1, ..., r\} \to [0, 1] \text{ such that} \\ \forall J \in m, \forall i \in \{1, 2, ..., r\}, u(t, i) = u(t', i) \quad = u(J, i) \quad \forall\ t, t' \in J \end{array} \right\}.$$
- the contrast is the log-likelihood defined by:
$$\gamma(s) \quad = \quad -\log\left[\prod_{t=1}^n \prod_{i=1}^r s(t, i)^{\mathbf{1}_{\{Y=i\}}}\right] = -\sum_{t=1}^n \sum_{i=1}^r \mathbf{1}_{\{Y=i\}} \log\left[s(t, i)\right].$$
- Associated to this contrast, the Kullback-Leibler information between s and u is:
$$K(s, u) = \sum_{t=1}^n \sum_{i=1}^r s(t, i) \log\left[\frac{s(t, i)}{u(t, i)}\right]$$
- For a given partition m, we obtain the minimum contrast estimator of s
$$\hat{s}_m(t, i) = \frac{N_J(i)}{\mid J \mid}, \text{ for } i \in \{1, ..., r\}, t \in J \text{ and } J \in m,$$

with $N_J(i) = \sum_{t\in J} \mathbf{1}_{\{Y=i\}}$, and the projection

$$\bar{s}_m(t,i) \quad = \quad \arg\min_{u\in\mathcal{S}_m} K(s,u) = \frac{\sum_{t\in J} s(t,i)}{\mid J \mid}$$

The following theorem gives the direct version of theorem 2.1.

Theorem 6.1. *Suppose that one observes independent variables Y_1, ..., Y_n taking their values in $\{1,2,...,r\}$ with $r \in \mathbb{N}$ and $r \geq 2$. We define for $t \in \{1,...n\}$ and $i \in \{1,2,...,r\}$*

$$P(Y_t = i) = s(t,i)$$

and consider a collection $\mathcal{M}_n$ of partitions constructed on the grid $\{1,...,n\}$. Let $(L_m)_{m\in\mathcal{M}_n}$ be some family of positive weights and define Σ as

$$\Sigma = \sum_{m\in\mathcal{M}_n} \exp(-L_m|m|) < +\infty.$$

Assume that

- *there exists some positive absolute constant ρ such that $s \geq \rho$,*
- *$\mathcal{M}_n$ is a collection of partitions constructed on a partition m_f such that $|J| \geq \Gamma[\log(n)]^2 \, \forall J \in m_f$ where Γ is a positive absolute constant.*

Let $\lambda_\varepsilon > 1/2$. If for every $m \in \mathcal{M}_n$

$$pen(m) \geq \lambda_\varepsilon \; r|m| \; \left(1 + 4\sqrt{L_m}\right)^2, \tag{9}$$

then

$$\mathbb{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] \leq C_{\lambda_\varepsilon} \inf_{m\in\mathcal{M}_n} \{K(s,\bar{s}_m) + pen(m)\} + C(\Sigma, r, \lambda_\varepsilon, \rho, \Gamma), \tag{10}$$

with $C_{\lambda_\varepsilon} = \frac{2(2\lambda_\varepsilon)^{1/3}}{(2\lambda_\varepsilon)^{1/3}-1}$.

We obtain this result by following the same lines of the proof of theorem 2.1. The main constant of interest is C_{λ_ε} which comes from the control of the term $\bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'})$ on a particular set $\Omega_{m_f}(\varepsilon)$. Here again this control is obtained using two results: the control of a chi-square statistic χ^2_m and the control $K(\bar{s}_m, \hat{s}_m)$ through a quantity V^2_m on $\Omega_{m_f}(\varepsilon)$.

6.2.2. *Control of $\bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'})$.* We write for $m' \in \mathcal{M}_n$

$$\bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'}) = \sum_{t=1}^{n}\sum_{i=1}^{r} \overline{\mathbf{1}_{\{Y_t=i\}}} \log\left[\frac{\hat{s}_{m'}(t,i)}{\bar{s}_{m'}(t,i)}\right].$$

By Cauchy-Schwarz inequality,

$$\bar{\gamma}(\bar{s}_{m'}) - \bar{\gamma}(\hat{s}_{m'}) \quad \leq \quad \sqrt{\sum_{i=1}^{r} \chi^2_{m'}(i)} \times V^2_{m'},$$

where

$$\chi^2_m = \sum_{i=1}^{r} \chi^2_{(m,i)} = \sum_{i=1}^{r}\sum_{J\in m} Z_J(i),$$

with $Z_J(i) = \frac{\overline{N_J(i)}^2}{\sum_{t\in J} s(t,i)}$ and $\overline{N_J(i)} = \sum_{t\in J}\left[\mathbf{1}_{\{Y=i\}} - s(t,i)\right]$, and

$$V^2_m \quad = \quad \sum_{J\in m}\sum_{i=1}^{r} |J|\, \bar{s}_m(J,i) \log^2\left[\frac{\bar{s}_m(J,i)}{\hat{s}_m(J,i)}\right].$$

Finally, the set $\Omega_{m_f}(\varepsilon)$ defined on a minimal partition m_f is defined in this context by

$$\Omega_{m_f}(\varepsilon) \quad = \quad \bigcap_{J\in m}\bigcap_{i=1}^{r} \left\{\left|N_J(i) - \sum_{t\in J} s(t,i)\right| \leq \varepsilon \sum_{t\in J} s(t,i)\right\}.$$

All those quantities are defined in the same manner than in our general approach, as we can recover $E_J^i = \sum_{t\in J} s(t,i)$ and $T_J^i = N_J(i)$. The main difference is that the dimension differs since $r = d+1$ terms are considered in each sum.

Control of $K(\bar{s}_m, \hat{s}_m)$ by V_m^2.

Lemma 6.2. *For all positive densities p and q with respect to μ, one has*

$$\frac{1}{2}\int f^2 \ \left(1 \wedge e^f\right) \ p \ d\mu \leq K(p,q) \leq \frac{1}{2}\int f^2 \ \left(1 \vee e^f\right) \ p \ d\mu$$

if one notes $f = \log\left(\frac{q}{p}\right)$.

Applying this lemma (for which a proof can be found in [15]), we have on $\Omega_{m_f}(\varepsilon)$

$$\begin{aligned} K(\bar{s}_m, \hat{s}_m) &\geq \frac{1}{2}\sum_{J\in m}\sum_{i=1}^{r} |J|\,\bar{s}_m(J,i)\left[1 \wedge \frac{\hat{s}_m(J,i)}{\bar{s}_m(J,i)}\right]\log^2\left[\frac{\bar{s}_m(J,i)}{\hat{s}_m(J,i)}\right] \\ &\geq \frac{1-\varepsilon}{2}V_m^2, \end{aligned}$$

and

$$\begin{aligned} K(\bar{s}_m, \hat{s}_m) &\leq \frac{1}{2}\sum_{J\in m}\sum_{i=1}^{r} |J|\,\bar{s}_m(J,i)\left[1 \vee \frac{\hat{s}_m(J,i)}{\bar{s}_m(J,i)}\right]\log^2\left[\frac{\bar{s}_m(J,i)}{\hat{s}_m(J,i)}\right] \\ &\leq \frac{1+\varepsilon}{2}\ V_m^2, \end{aligned}$$

resulting in

$$\frac{1-\varepsilon}{2}V_m^2 \leq K(\bar{s}_m, \hat{s}_m) \leq \frac{1+\varepsilon}{2}V_m^2, \tag{11}$$

Exponential bound for χ_m^2. We want to control the chi-square statistic around its expectation by using Bernstein's inequality as in Section 3. Here the required control of $N_J(i)$ can be obtained through a direct application of the bounded version of Bernstein since $\mathbf{1}_{\{Y=i\}}$ is bounded by 1. We get:

$$P\left[\left|\sum_{t\in J}\overline{\mathbf{1}_{\{Y=i\}}}\right| \geq x\right] \quad \leq \quad 2\exp\left(-\frac{x^2}{2\left(\frac{x}{3} + \sum_{t\in J} s(t,i)\right)}\right),$$

which, as in the general approach, can be used to obtain

$$P\left(\Omega_{m_f}(\varepsilon)^c\right) \leq \frac{C(\Gamma, \rho, a, r, \varepsilon)}{n^a}.$$

Then, since

$$\mathbb{E}\left[\chi^2(m,i)\right] = \sum_{J\in m}\frac{Var\left[\sum_{t\in J}\mathbf{1}_{\{Y=i\}}\right]}{\sum_{t\in J} s(t,i)} = \sum_{J\in m}\frac{\sum_{t\in J} s(t,i)\left[1-s(t,i)\right]}{\sum_{t\in J} s(t,i)}.$$

and $\forall t, i\ \rho \leq s(t,i) \leq 1$, and $r-1 \geq r/2$ for $r \geq 2$, we have the following bounds

$$\rho\frac{r}{2}|m| \leq \sum_{i=1}^{r}\mathbb{E}_s\left[\chi^2(m,i)\right] \leq r|m|. \tag{12}$$

By controlling the moment of $Z_J(i)$ (as in Section 3) and using Bernstein's inequality again, we conclude to

$$P\left(\sum_{i=1}^{r}\chi^2(m,i)\mathbf{1}_{\Omega_{m_f}(\varepsilon)} \ \geq r|m| + 8r\left(1+\frac{\varepsilon}{3}\right)\sqrt{x|m|} + 4r\left(1+\frac{\varepsilon}{3}\right)x\right) \quad \leq \quad r\exp(-x).$$

Control of $K(\bar{s}_m, \hat{s}_m)$ by χ^2_m. The term V_m^2 can be written as follows:

$$V_m^2 = \sum_{J\in m}\sum_{i=1}^{r} |J| \frac{\left[\hat{s}_m(J,i) - \bar{s}_m(J,i)\right]^2}{\bar{s}_m(J,i)} \left[\frac{\log\left[\hat{s}_m(J,i)/\bar{s}_m(J,i)\right]}{\hat{s}_m(J,i)/\bar{s}_m(J,i) - 1}\right]^2,$$

so that using $\frac{1}{1\vee x} \leq \frac{\log x}{x-1} \leq \frac{1}{1\wedge x}$ for all $x > 0$, on the set $\Omega_{m_f}(\varepsilon)$, we have:

$$\frac{1}{(1+\varepsilon)^2}\sum_{i=1}^{r}\chi^2(m,i) \leq V_m^2 \leq \frac{1}{(1-\varepsilon)^2}\sum_{i=1}^{r}\chi^2(m,i).$$

Combining this equation with relation (11) gives, on $\Omega_{m_f}(\varepsilon)$

$$\frac{1-\varepsilon}{2(1+\varepsilon)^2}\sum_{i=1}^{r}\chi^2(m,i) \leq K(\bar{s}_m, \hat{s}_m) \leq \frac{1+\varepsilon}{2(1-\varepsilon)^2}\sum_{i=1}^{r}\chi^2(m,i). \tag{13}$$

6.2.3. *Control of the two other terms.* We explicit here the control of $\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)$ and $\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})$ which appeared in the negligible constant of the risk.

- Control of the term $\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)$. We have to control

$$\mathbb{E}_s\left[\left(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)\right)\mathbf{1}_{\Omega_{m_f}(\varepsilon)(\varepsilon)}\right] = -\mathbb{E}_s\left[\left(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)\right)\mathbf{1}_{\Omega_m(\varepsilon)^c}\right],$$

  and we bound this expectation by

$$\begin{aligned}\left|\mathbb{E}\left[\left(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)\right)\mathbf{1}_{\Omega_{m_f}(\varepsilon)}\right]\right| &\leq \left|\mathbb{E}\left[\left(\bar{\gamma}(\bar{s}_m) - \bar{\gamma}(s)\right)\mathbf{1}_{\Omega_{m_f}(\varepsilon)^c(\varepsilon)}\right]\right| \\ &\leq nr\log\left(\frac{1}{\rho}\right)P\left(\Omega_{m_f}(\varepsilon)^c\right) \\ &\leq \frac{C(\Gamma,\rho,a,r,\varepsilon)}{n^{a-1}}\end{aligned}$$

- Control of term $\bar{\gamma}(\hat{s}_{m'}) - \bar{\gamma}(\bar{s}_{m'})$. We use the following proposition

  Proposition 6.3. *For every $u \in \mathcal{S}$ and any positive x,*

$$P\left(\bar{\gamma}(s) - \bar{\gamma}(u) \geq K(s,u) - 2h^2(s,u) + 2xr\right) \leq r\exp(-x),$$

  where

$$\bar{\gamma}(s) - \bar{\gamma}(u) = \sum_{i=1}^{r}\sum_{t=1}^{n}\overline{\mathbf{1}_{\{Y_t=i\}}}\log\left[\frac{u(t,i)}{s(t,i)}\right],$$

  and h^2 is the Hellinger distance.

  The proof relies on the same arguments than for proposition 4.3.

6.2.4. *Proof of theorem 6.1.* The proof of this theorem is obtained by following the lines of Section 4.2.4: we introduce sets of large probability, $\Omega_1(\xi)$ and $\Omega_2(\xi)$ and gather the previous results on the control of each terms of the main decomposition. This results in the main inequality :

$$\mathbb{E}\left[h^2(s,\hat{s}_{\hat{m}})\right] \leq \frac{2(2\lambda_\varepsilon)^{1/3}}{(2\lambda_\varepsilon)^{1/3} - 1}\left\{K(s,\bar{s}_m) + pen(m)\right\} + C(\Sigma, r, \lambda_\varepsilon, \rho, \Gamma).$$

6.2.5. *Oracle-type inequality.* The following proposition gives the risk of an estimator $\hat{s}_m$ for any $m \in \mathcal{M}_n$.

Proposition 6.4. *Under the assumptions of theorem 6.1, we have*

$$\mathbb{E}\left[K(s,\widehat{s}_m)\right] \geq K(s,\bar{s}_m) + \frac{1-\varepsilon}{4(1+\varepsilon)^2}\ \rho\ r\ |m| - \frac{C(\Gamma,\rho,a,r,\varepsilon)}{n^{a-1}}, \tag{14}$$

where $a > 2$ and $C(\Gamma,\rho,a,r,\varepsilon)$ is a positive constant only depending on Γ, a, ρ, r and ε.

Proof: According to inequalities (12) and (13), we get the following lower bound of the risk:

$$\begin{aligned}\mathbb{E}\left[K(s,\hat{s}_m)\right] &= K\left(s,\bar{s}_m\right)+\mathbb{E}\left[K(\bar{s}_m,\hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)}\right]+\mathbb{E}\left[K(\bar{s}_m,\hat{s}_m)\mathbf{1}_{\Omega_{m_f}(\varepsilon)^c}\right]\\ &\geq K\left(s,\bar{s}_m\right)+\frac{1-\varepsilon}{4\left(1+\varepsilon\right)^2}\rho r|m|-\frac{C\left(\Gamma,\rho,a,r,\varepsilon\right)}{n^{a-1}}.\end{aligned}$$

Using the results of proposition 6.4 and theorem 6.1, we obtain the following oracle-type inequality:

Corollary 6.5. *Let pen : $\mathcal{M}_n \to \mathbb{R}_+$ be defined by (9), and assume $s \geq \rho$ for some $\rho > 0$. Then there exists some constant C such that for an exhaustive search,*

$$\mathbb{E}\left[h^2\left(s,\hat{s}_{\hat{m}}\right)\right] \leq C\log\left(n\right)\inf_{m\in\mathcal{M}_n}\left\{K\left(s,\hat{s}_m\right)\right\}+C(\Gamma,\rho,a,r,\lambda_\varepsilon,\Sigma). \tag{15}$$

Remark: $K\left(s,\bar{s}_m\right)$ is close to $4h^2\left(s,\bar{s}_m\right)$ when $||\log(s/\bar{s}_m)||_\infty < \infty$ and under this assumption the above upper bound can be stated as follows

$$\mathbb{E}\left[h^2\left(s,\hat{s}_{\hat{m}}\right)\right] \leq C\log\left(n\right)\inf_{m\in\mathcal{M}_n}\mathbb{E}\left[h^2\left(s,\hat{s}_m\right)\right]+C(\Gamma,\rho,a,r,\lambda_\varepsilon,\Sigma).$$

6.3. **Comparison of the constants.** First, we aim at comparing the bounds of the risk of our estimator between the general and direct approach (see equations (1) and (10)). In the general approach, the constant is expressed as $C_{\beta_\varepsilon} = \frac{1+C_2^{-1}(\varepsilon)}{C_2^{-1}(\varepsilon)}$ where $C_2(\varepsilon) = \beta_\varepsilon = \frac{E^{max}\kappa}{2m_\varepsilon}(1+\varepsilon)^3$. Using the results of Section 6.1, we have $E^{max} < 1$, $\kappa = 2$ and $m_\varepsilon = (E^{min}-\varepsilon)^{d+1}$, so that $C_2(\varepsilon)$ can be bounded by

$$C_2(\varepsilon) \leq \frac{(1+\varepsilon)^3}{(E^{min}-\varepsilon)^r} \leq \left(\frac{1+\varepsilon}{E^{min}-\varepsilon}\right)^\alpha \quad \text{where } \alpha = \max\{r,3\}.$$

We therefore obtain:

$$C_{\beta_\varepsilon} \leq \frac{(1+E^{min})\beta_\varepsilon^{1/\alpha}}{E^{min}\beta_\varepsilon^{1/\alpha}-1}$$

which can further be bounded by $C_{\beta_\varepsilon} = \dfrac{2\beta_\varepsilon^{1/\alpha}}{\frac{\beta_\varepsilon^{1/\alpha}}{a+a^2d}-1}$ with $\beta_\varepsilon > (a+a^2d)^\alpha$.
In the direct approach, the constant is $C_{\lambda_\varepsilon} = \frac{2(2\lambda_\varepsilon)^{1/3}}{(2\lambda_\varepsilon)^{1/3}-1}$ with $\lambda_\varepsilon = \frac{1}{2}\frac{(1+\varepsilon)^3}{(1-\varepsilon)^3} > \frac{1}{2}$.

In both cases, β_ε and λ_ε are the constants from the penalty function that are tuned from the data. We can notice that the general shape of C_{β_ε} and C_{λ_ε} are the same, except that the power of the penalty constant is directly related to the number of categories in the general case while it is fixed to 3 in the direct approach.
This is compensated by, on one hand, a greater constraint on β_ε since $a > 1$ and thus $\beta_\varepsilon > r^r$ while $\lambda_\varepsilon > 1/2$, and on the other hand by a different multiplicative factor in the penalty function, since we obtain $pen(m) > \beta_\varepsilon d|m|(1+4\sqrt{L_m})^2$ in the general approach and $pen(m) > \lambda_\varepsilon(d+1)|m|(1+4\sqrt{L_m})^2$ in the direct approach.

Then comparing the oracle inequalities given by (4) and (15) in the general and direct cases respectively results in comparing the constants C that is:

- $C_{\beta_\varepsilon}\frac{\beta_\varepsilon}{C(\varepsilon)}$ where $C(\varepsilon)$ (in equation (3)) is $C(\varepsilon) = \frac{1}{2}\zeta E^{min}\dfrac{m_\varepsilon}{M_\varepsilon}$ in the general approach, *i.e.* $C(\varepsilon) = \frac{1}{2}E^{min}(1-E^{max})m_\varepsilon$
- $C_{\lambda_\varepsilon}\frac{\lambda_\varepsilon}{C(\varepsilon)}$ where $C(\varepsilon) = \frac{1-\varepsilon}{4(1+\varepsilon)^2}\rho$ (cf equation (14)) in the direct one.

$C_{\beta_\varepsilon}\beta_\varepsilon$ and $\lambda_\varepsilon C_{\lambda_\varepsilon}$ behaving as previously, one looses at least a constant $\frac{1}{2\zeta(E^{min})^{d+1}}$ between the general and the direct approaches.

7. Simulation study and application to DNA sequences

In this section, we apply our estimator in two scenarios. The first is a simulation study where the observations are sampled from the exponential distribution with piece-wise constant rate parameter. In this case the distribution is continuous and the sufficient statistic is unidimensional (and is the variable itself). The second is an application to a real data-set on the analysis of DNA sequences in terms of base-composition. The objective is to apply a segmentation model in order to find homogeneous regions which can be related to structural and functional biological regions. In this case we model the observations with a categorical distribution with piece-wise constant proportion parameters. The distribution is therefore discrete and the sufficient statistic is multivariate (equal to the indicator of the variable taking a category value).

7.1. **Simulation study with exponential distribution.** We simulated 400 datasets of length $n = 10^5$ for which the number of segments K was drawn from a Poisson distribution with mean $\bar{K} = 50$, and the $K - 1$ change-points were sampled uniformly on $\{2, \dots, n-1\}$ subject to the constraint that segments had to be of length at least 10. We considered 4 sets of values for the rate parameters. In all scenarios, odd segments had a rate of 0.01 while the rate on even segments was chosen randomly with probability 0.4, 0.2, 0.3 and 0.1 among the values $(0.05, 0.1, 0.02, 0.005)$ for datasets 1 through 100, $(0.02, 0.05, 0.015, 0.005)$ for datasets 101 through 200, $(0.015, 0.02, 0.0125, 0.007)$ for datasets 201 through 300, and $(0.015, 0.02, 0.011, 0.008)$ for datasets 301 through 400. This resulted in datasets similar to those shown in Figure 1.

In this framework, the segmentation was performed using the pruned dynamic programming algorithm [17, 18] from which we obtained the optimal segmentation for every K up to $Kmax = 200$. The number of segments was then obtained using the penalty function proposed in (2) where the constant is calibrated using the slope heuristic (see [13]).

To assess the quality of our results, we introduce, as in [19], the quantities $\mathcal{E}(m_{\hat{s}}||m_s)$ and $\mathcal{E}(m_s||m_{\hat{s}})$ between the partitions associated with the true and estimated distributions s and $\hat{s}$, where

$$\mathcal{E}(m_A||m_B) = \sup_{b \in m_B} \inf_{a \in m_A} |a - b|.$$

The first quantity assesses how our estimated segmentation is able to recover the true change-points. Intuitively, the segmentation with the largest number of segments will have the greatest chance of yielding a small value of $\mathcal{E}(m_{\hat{s}}||m_s)$. On the contrary, the second quantity judges how relevant the proposed change-points are compared to the true partition: a segmentation with too many segments will necessarily have change-points far from the true ones. Note that the Hausdorff distance can then be recovered as $\sup\{\mathcal{E}(m_{\hat{s}}||m_s), \mathcal{E}(m_s||m_{\hat{s}})\}$.

The performance of our procedure is finally assessed *via*:

- The difference between the true number of segments and the estimated, $\Delta = K - \hat{K}$;
- $\mathcal{E}(m_{\hat{s}}||m_s)$ and $\mathcal{E}(m_s||m_{\hat{s}})$, between our estimator and the true segmentation;
- The Kullback distance between true and estimated distributions, namely $\sum_t \left(\frac{\hat{\lambda_t}}{\lambda_t} - \log \frac{\hat{\lambda_t}}{\lambda_t} - 1 \right)$; and
- The Hellinger distance between true and estimated distributions, namely $\sum_t \left(1 - 2\frac{\sqrt{\lambda_t \hat{\lambda}_t}}{\lambda_t + \hat{\lambda}_t} \right)$.

To allow a fair assessment of our estimator, the Haussdorff, Kullback and Hellinger criteria are also computed for the optimal segmentation for the true number of segments, which we will denote $\hat{m}_K$, and the resulting estimated distribution $\hat{s}_{\hat{m}_K}$.

Our method leads to a tendency to under-estimate the number of segments, as is shown in Figure 2 a). As in classical studies of model selection for segmentation, our estimator tends to under-estimate the number of segments in particular when the scenario becomes more difficult to segment. Indeed, the number of segments is reduced in order to avoid false detection. This phenomenon is illustrated in Figure 2 b) as our estimator (represented through the blue boxplots) yields high values of $\mathcal{E}(m_{\hat{s}}||m_s)$ due to

the missed segments (subfigure b_1) but low values of $\mathcal{E}(m_s||m_{\hat{s}})$ as the segments we propose tend to correspond to true segments (subfigure b_2). On the contrary, the segmentation $\hat{s}_{\hat{m}_K}$ corresponding to the

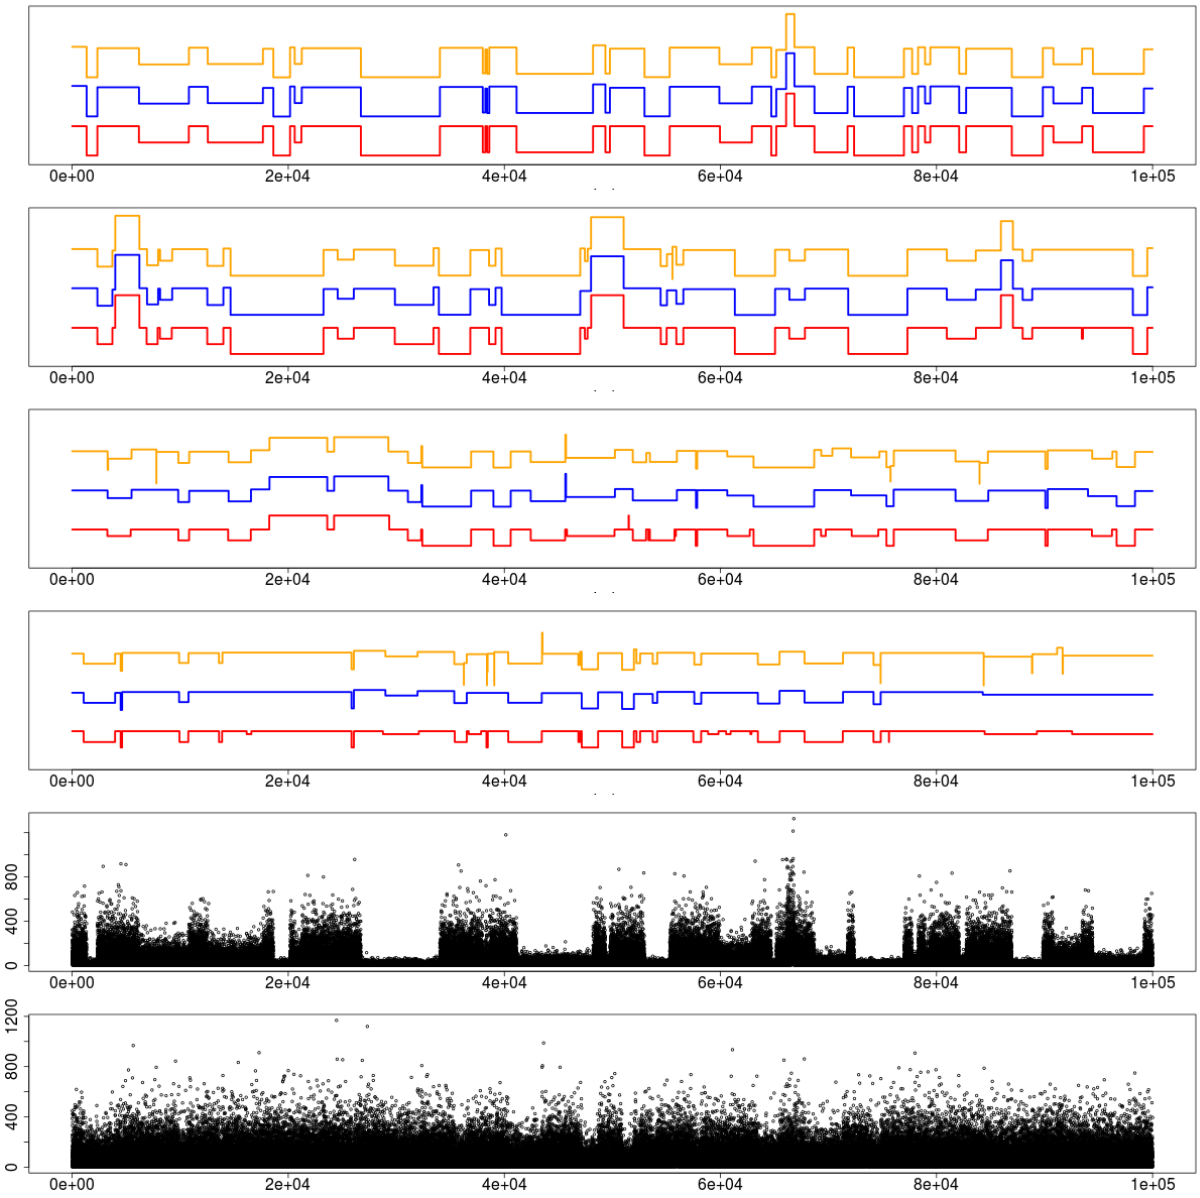

FIGURE 1. **Simulated datasets and segmentation.** The four first figures show examples of simulation design for each of the four groups of rate values; the last figure shows an example of dataset to segment in the last group of values. Red lines (bottom) indicate the true distribution (the inverse of the rate, $1/\lambda$, is represented) while blue lines (middle) indicate the estimated distribution and the orange lines (top) indicate the optimal segmentation (w.r.t. likelihood loss) for the true number of segments.

true number of segments (yellow boxplots) has lower values of $\mathcal{E}(m_{\hat{s}}||m_s)$ as it has more change-points thus lower distances to the missed ones, but higher values of $\mathcal{E}(m_s||m_{\hat{s}})$ as some of its segments are spurious. On a particular example, presented in Figure 1 by comparing the red line (true segmentation) to the blue one (corresponding to our estimator), and more so on the bottom subfigure illustrating a simulation from the fourth group, we observe that our method tends to fail to recover small segments with rate very close to surrounding ones, whereas the optimal segmentation (as represented by the yellow line) still fails

to recover those small segments, thus proposing additional ones, typically very short (average length=9) with very different rate values.
Finally, Figures 2 c) and d) show that in terms of Kullback-Leibler and Hellinger distances, our estimator performs at least as well than the one with the true number of segments, and significantly better when the profiles become more difficult to segment.

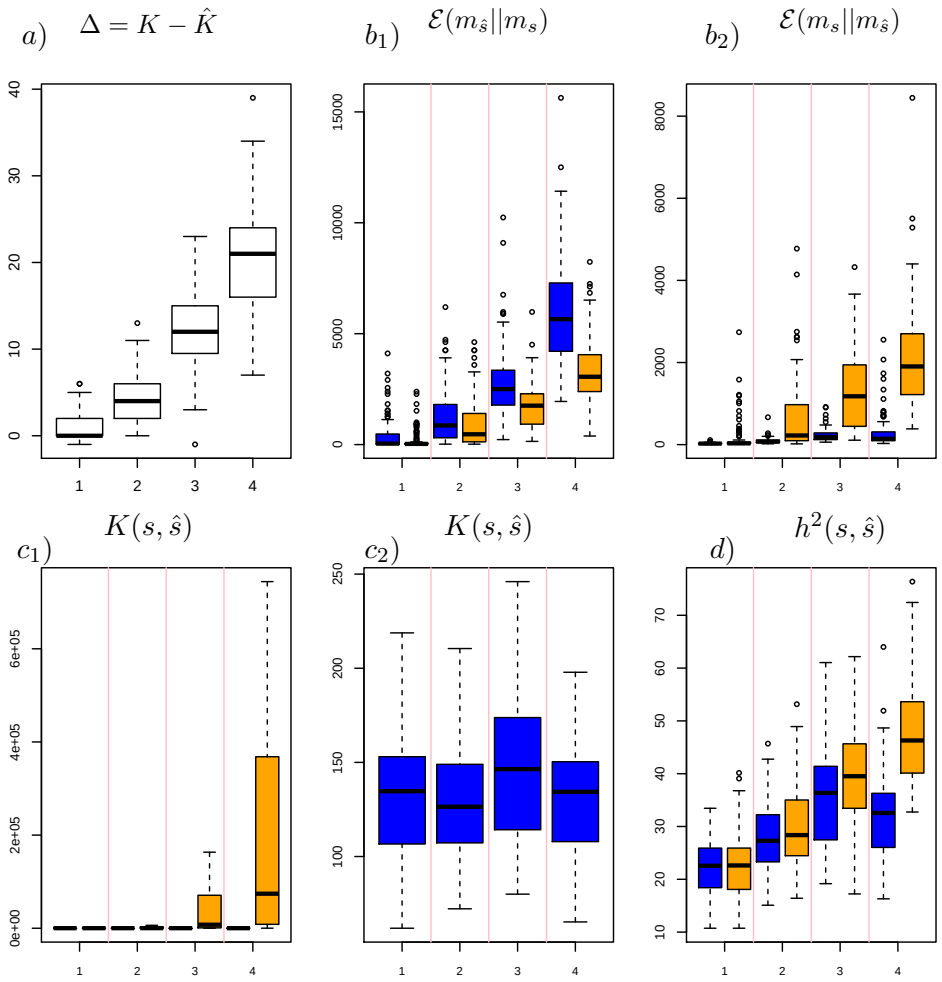

FIGURE 2. **Evaluation of the performance criteria in the four simulation study frameworks.** a) difference between the true number of segments and the estimated. Our method tends to underestimate the number of segments, with performance degradating as the segmentation difficulty increases. $b_1)$ boxplot of the values of $\mathcal{E}(m_{\hat{s}}||m_s)$ and $b_2)$ $\mathcal{E}(m_s||m_{\hat{s}})$ over the hundred simulations in each framework. In each case, left boxplots (blue) assess our estimator, right boxplots (orange) assess the estimator corresponding to the true number of segments. $c_1)$ and $c_2)$ Boxplot of the Kullback-Leibler distance to the true distribution for the estimated distribution (blue), and the optimal distribution in K segments (orange) in each simulation framework. In $c_2)$, only $\hat{s}_{\hat{K}}$ is shown. $d)$ Same as $c_1)$ but for the Hellinger distance.

7.2. **Application of a DNA sequence.** The objective of this application is to find regions of a DNA sequence which are homogeneous in terms of base composition, that is which present a stability in the frequencies of the four nucleotide letters. These regions are thought to correspond to areas of the genome which are biologically significant. To this end, we apply our procedure modeling the data with categorical variables with $d = 3$ (see Section 6 for the model).

Here we consider the bacteriophage Lambda genome with length $n = 48502$ base pairs which is a parasite of the intestinal bacterium *Escherichia coli*. This genome has been used for the comparison of segmentation methods (see [20] and references therein) such as HMM ([21], [22]) or penalized quasi-likelihood [23]. The data and its annotation are publicly available from the National Center for Biotechnology Information (NCBI) pages at the url adress http://www.ncbi.nlm.nih.gov/.
From a computational point of view, the large size of the Lambda genome hampers the direct use of the classical Dynamic Programming (DP) algorithm. Here we propose a hybrid algorithm that consists in first selecting a small number of relevant change-points using the CART algorithm [24], and then using dynamic programming on this set of candidate change-points. As in the simulation study, the penalty constant is calibrated using the slope heuristic (see [13]). Four change-points (*i.e.* five segments) are selected by our criterion at positions 22546, 27829, 38004 and 46528. The associated regions are characterized by different base composition as shown in Figure 3 which represents the estimated probabilities of each base for the obtained segmentation.

These change-points are very close (and even on some occasion the precise same) as the one obtained in [23], which concluded to 3 more change-points. This reference also supposes bases to be independent and uses a penalized contrast procedure to perform the segmentation, and is in this sense the closest approach to ours. The segments we identify reflect changes in transcription direction. Indeed, this direction is forward up to base 20855, it then switches to reverse from base 22686 to 37940, switches back to forward between 38041 to 46427 and finally reverse again from 46459 to the end. Note that a refinement has been obtained when assuming a dependence relationship between bases (see [21] and [22]).

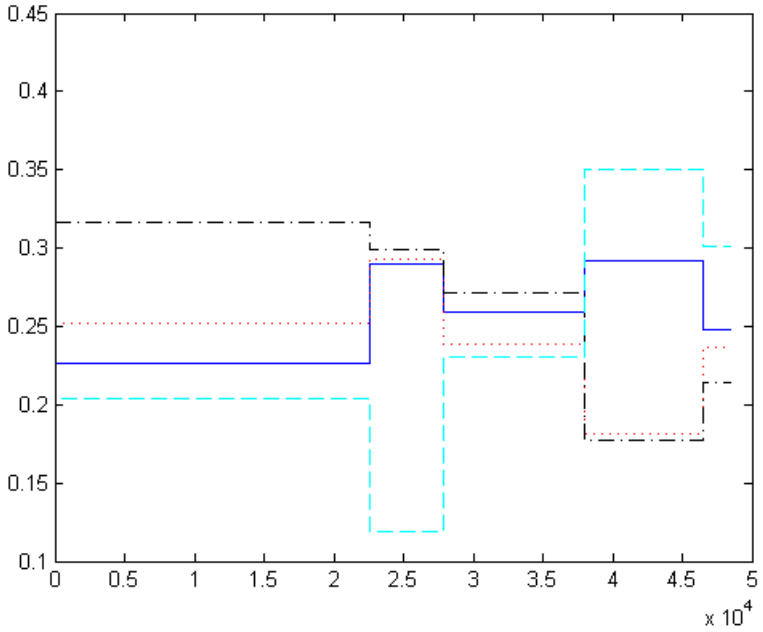

FIGURE 3. Estimated probabilities on each segment of the selected segmentation: blue and '-' for the base adenosine A, red and ':' for the base cytosine C, black and '-.' for the base guanine G and cyan and '–' for the base thymine T.

8. CONCLUSION

We have proposed a general approach to the selection of the number of segments in the general framework where the data can be modeled using a distribution from the exponential family. As expected, the log-partition function and its many properties are instrumental in the derivation of the bounds and the obtention of the oracle inequality.

The main drawback of our approach is that it can only be applied to distributions for which the sufficient statistics can be guaranteed to belong to a positive set. This is mainly due to the use of the chi-square statistic which was initially defined for the analysis of count (and thus positive) data. It is very likely that our result could be extended to overcome this issue by modifying decomposition 5 and controlling the fluctuation of the statistics around their expectation using different concentration inequalities. However, the main goal of this work, beside providing a general penalty function for model selection, was to underline the role of the log-partition function in our previous work. Similarly, this work should easily be extended to multivariate distributions from the exponential family.
Here, using the particular case of categorical variables as a example, we have shown that the loss in tightness of the main constant is not a drastic issue by comparing the results obtained from the general approach to that from the direct one.

Moreover, we have shown in many examples through simulation and application studies (negative binomial and Poisson distributions in our previous work, exponential distribution in the simulation study and categorical distribution in the application to DNA sequences) that our approach is a powerful method to detect significant changes in the distribution of the data, which can often be related to phenomenon of interest. It outperforms existing criteria (see for instance [2]) with a behavior that is expected in classical studies of segmentation problems: it tends to under-estimate the number of segments in order to avoid false detection.

Acknowledgements The authors would like to thank Elodie Nédélec for her help in the categorical study, and Mahendra Mariadassou and Hoel Queffelec for helpfull discussions on exponential families and their geometric properties.

Appendix A. Computations for classic laws of the exponential family

Because the constants are strongly related to properties of the cumulants, we distinguish distributions with explicit cumulants from those that can be computed through some recurrence property.

We first study laws for which the cumulants are given explicitly. They include the Poisson, exponential, Gaussian distributions, etc...

- In the case of the Poisson distribution with parameter λ, the log-partition function is analytic on $\mathbb{R}$ and all cumulants are equal to e^θ so we can use $\zeta = \kappa = 1$.

- In the Gaussian case $\mathcal{N}(\mu, \sigma^2)$, we require the mean parameter μ to be strictly positive in order to define the χ^2 statistic. In this case, the natural parameters are given by $\boldsymbol{\theta} = \left(\frac{\mu}{\sigma^2}, -\frac{1}{2\sigma^2}\right)$, and the sufficient statistic is $T(Y) = (Y, Y^2)$, for which the cumulants are given by explicit formulaes through the log-partition fuction $A(\boldsymbol{\theta}) = -\frac{\theta_1^2}{4\theta_2} - \frac{1}{2}\log(\theta_2)$.
  In particular, the cumulants of Y are simply $c_1 = -\frac{\theta_1}{2\theta_2}$, $c_2 = -\frac{1}{2\theta_2}$ and $c_k = 0$ for $k \geq 3$, and we have $Var(Y) = \theta_1^{-1}\mathbb{E}(Y)$ so that $\zeta_1 = 1/\theta_1^{max}$ and $\kappa_1 = 1/\theta_1^{min}$ work.
  For the second order statistic Y^2, for $k \geq 2$ the cumulants are given by $c_k = \left(-\frac{1}{4}\frac{\theta_1^2}{\theta_2} + \frac{1}{2k}\right)\left(-\frac{1}{\theta_2}\right)^k k!$ on $-|\theta_2| < z < |\theta_2|$ and we have $-\frac{1}{\theta_2}\mathbb{E}(Y^2) \leq Var(Y^2) \leq -\frac{2}{\theta_2}\mathbb{E}(Y^2)$. With the previous notations, we can choose $\tau = -1/\theta_2^{min}$ and $\zeta_2 = -1/\theta_2^{min}$ and $\kappa_2 = -2/\theta_2^{max}$ work.

- The case of the Pareto distribution also has to be reduced to a fixed scale parameter x_m and a shape parameter α smaller than $-1/\log x_m$. We have $\zeta = \dfrac{1}{\theta^{max}}\dfrac{1}{\theta^{max}\log x_m - 1}$. The cumulants are then simply obtained as $c_k = (-1)^k\frac{(k-1)!}{\theta^k}$ for $k > 1$ and we can take $\kappa = \frac{\log x_m}{\theta^{min}\log x_m - 1}$.

- The Gamma distribution $\Gamma(\alpha, \beta)$ with known shape parameter α has an analytic log-partition function with radius $|\theta|$ where θ is the opposite of the scale parameter β. The sufficient statistic

Y has cumulants given by $c_k = (-1)^k(k-1)!\alpha\theta^{-k}$ and therefore τ can be chosen as $-\frac{1}{\theta^{max}}$ and finally $\zeta = -1/\theta^{min}$ and $\kappa = -1/\theta^{max}$.
The inverse Gamma distribution with known shape parameter yields the same result since its sufficient statistic is $1/Y$ which, by definition, follows a Gamma distribution.

- The centered Weibull distribution with shape parameter k and scale parameter λ belongs to the exponential family provided the shape parameter k is known. The sufficient statistic is then Y^k, which, when normalized by λ, follows an exponential distribution with parameter 1. Results on ζ and κ can therefore easily be deducted from the exponential study.

- The Laplace distribution with known mean positive parameter μ (positivity is assumed for the definition of χ^2) and scale parameter b has sufficient statistic $|Y-\mu|$ and natural parameter $\theta = -1/b$. As the sufficient statistic follows an exponential distribution with parameter $-\theta$, all results are easily deducted from this study.

We then study distributions for which cumulants can be obtained by some recurrence property. This includes the Bernoulli, binomial (with known parameter n), negative binomial (with known overdispersion parameter ϕ), etc... While the recurrence are not derived properly in this section, a complete example is detailed in the application to categorical variables in Section 6.

- The Bernoulli and binomial distributions with fixed number of trials n can be treated together by considering $n = 1$ in the first case. The sufficient statistic is Y and the natural parameter is $\theta = \log\frac{p}{1-p}$. The log-partition function, $A(\theta) = n\log\left(1+e^{\theta}\right)$ is analytic with radius of convergence $\log\left(1+\frac{1+e^{\theta}}{e^{\theta}}\right)$ and we can take $R = \log 2$. While the cumulants are not given by an explicit formula, they can be obtained with the recurrence : $c_1 = np$ and $c_{k+1} = p(1-p)\frac{\delta c_k}{\delta p}$ for $k \geq 1$, which is the same as in the categorical study. Similar computations lead to the results of Table 1.

- Similarly, the negative binomial distribution belongs to the exponential family provided the overdispersion parameter ϕ is known. The natural parameter is then $\theta = \log(1-p)$ and the sufficient statistic is Y. Cumulants verify the recurrence property $c_1 = \phi\frac{1-p}{p}$ and $c_{k+1} = (1-p)\frac{\delta c_k}{\delta(1-p)}$ and we can easily show that $c_k = \phi\frac{1-p}{p^k}f_k(p)$ with $f_k(p)$ a polynomial in p with degree at most $k-2$ and sum of absolute coefficients less than $2^{k-3}k!$. This leads to the choices $\zeta = 1$ and $\kappa = e^{-\theta_{min}}$.

References

[1] Lebarbier E: **Detecting multiple change-points in the mean of Gaussian process by model selection**. *Signal Processing* 2005, **85**(4):717–736.

[2] Cleynen A, Lebarbier E: **Segmentation of the Poisson and negative binomial rate models: a penalized estimator**. *ESAIM: Probability and Statistics* 2014.

[3] Akaike H: **Information Theory and Extension of the Maximum Likelihood Principle**. *Second international symposium on information theory* 1973, :267–281.

[4] Yao YC: **Estimating the number of change-points via Schwarz' criterion**. *Statistics & Probability Letters* 1988, **6**(3):181–189.

[5] Rigaill G, Lebarbier E, Robin S: **Exact posterior distributions and model selection criteria for multiple change-point detection problems**. *Statistics and Computing* 2012, **22**(4):917–929.

[6] Birgé L, Massart P: **Minimal penalties for Gaussian model selection**. *Probability Theory Related Fields* 2007, **138**(1-2):33–73.

[7] Zhang NR, Siegmund DO: **A modified Bayes information criterion with applications to the analysis of comparative genomic hybridization data**. *Biometrics* 2007, **63**:22–32.

[8] Brown LD: **Fundamentals of statistical exponential families with applications in statistical decision theory**. *Lecture Notes-monograph series* 1986, :i–279.

[9] Wainwright MJ, Jordan MI: **Graphical models, exponential families, and variational inference**. *Foundations and Trends® in Machine Learning* 2008, **1**(1-2):1–305.

[10] Kakade SM, Shamir O, Sridharan K, Tewari A: **Learning exponential families in high-dimensions: Strong convexity and sparsity**. *arXiv preprint arXiv:0911.0054* 2009.

[11] Lee JD, Sun Y, Taylor J: **On model selection consistency of M-estimators with geometrically decomposable penalties**. *arXiv preprint arXiv:1305.7477* 2013.
[12] Massart P: *Concentration inequalities and model selection*. Springer Verlag 2007.
[13] Arlot S, Massart P: **Data-driven calibration of penalties for least-squares regression**. *The Journal of Machine Learning Research* 2009, **10**:245–279.
[14] Birgé L, Massart P: **Gaussian model selection**. *Journal of the European Mathematical Society* 2001, **3**(3):203–268.
[15] Castellan G: **Modified akaike's criterion for histogram density estimation**. *C. R. Acad. Sci., Paris, Sér. I, Math. 330* 2000, **8**:729–732.
[16] Baraud Y, Birgé L: **Estimating the intensity of a random measure by histogram type estimators**. *Probability Theory Related Fields* 2009, **143**(1-2):239–284.
[17] Rigaill G: **Pruned dynamic programming for optimal multiple change-point detection**. *Arxiv:1004.0887* 2010, [[http://arxiv.org/abs/1004.0887]].
[18] Cleynen A, Koskas M, Lebarbier E, Rigaill G, Robin S: **Segmentor3IsBack: an R package for the fast and exact segmentation of Seq-data.** *Algorithms for Molecular Biology* 2014, **9**:6.
[19] Harchaoui Z, Lévy-Leduc C: **Multiple change-point estimation with a total variation penalty**. *Journal of the American Statistical Association* 2010, **105**(492).
[20] Braun JV, Muller HG: **Statistical methods for DNA sequence segmentation**. *Statistical Science* 1998, :142–162.
[21] Boys RJ, Henderson DA: **A bayseian approach to DNA Sequence segmentation**. *Biometrics* 2004, **60**(2):573–588.
[22] Muri F: **Modelling bacterial genomes using hidden Markov models**. *Compstat98. Proceedings in Computational Statistics, Eds R. Payne and P. Green* 1998, :89–100.
[23] Braun JV, Braun R, Müller HG: **Multiple changepoint fitting via quasilikelihood, with application to DNA sequence segmentation**. *Biometrika* 2000, **87**(2):301–314.
[24] Breiman, Friedman, Olshen, Stone: **Classification and Regression Trees.** *Wadsworth and Brooks* 1984.

Harvard School of Public Health, Boston, USA
E-mail address: acleynen@jimmy.harvard.edu

AgroParisTech UMR518, Paris 5e, France

INRA UMR518, Paris 5e, France.
E-mail address: emilie.lebarbier@agroparistech.fr